AD-A104 369 APPLIED DECISION ANALYSIS INC MENLO PARK CA F/G 15/3
DECISION AIDS FOR NAVY TACTICAL ANTI-AIR WARFARE.(U)
JUN 81 A C MILLER, L J POLLENZ, F SHIH N00014-80-C-0405

UNCLASSIFIED NL

1 of 2
AD
A104369

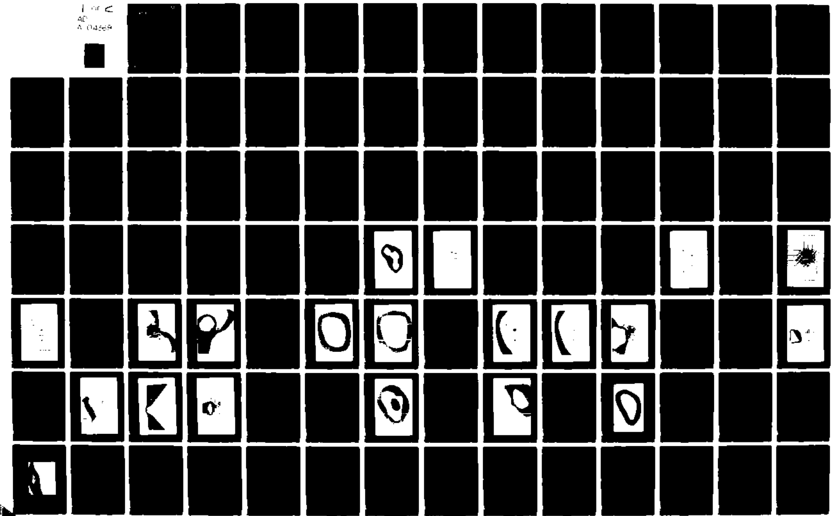

LEVEL II

12

APPLIED DECISION ANALYSIS, inc.

3000 Sand Hill Road, Menlo Park, CA 94025 (415) 854-7101

AD A104369

June 1981

DECISION AIDS FOR NAVY TACTICAL ANTI-AIR WARFARE

Prepared for:

Dr. Martin A. Tolcott
Director
Engineering Psychology Program
Code 455
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Prepared by:

Dr. Allen C. Miller
Dr. Lynne J. Pollenz
Mr. Frank Shih, Jr.
Applied Decision Analysis, Inc.
3000 Sand Hill Road
Menlo Park, CA 94025

"Original contains color plates: All DTIC reproductions will be in black and white"

81 9 17 040

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER / 2. GOVT ACCESSION NO. AD-A104 369	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Decision Aids For Navy Tactical Anti-Air Warfare	5. TYPE OF REPORT & PERIOD COVERED Final Rep[redacted]t,
	6. PERFORMING ORG. REPORT NUMBER ADA Project #2026
7. AUTHOR(s) Allen C. Miller, III; Lynne J. Pollenz; Frank Shih, Jr.	8. CONTRACT OR GRANT NUMBER(s) N00014-80-C-0405
9. PERFORMING ORGANIZATION NAME AND ADDRESS Applied Decision Analysis, Inc. 3000 Sand Hill Road Menlo Park, CA 94025	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217	12. REPORT DATE Jun[redacted]981
	13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Defense Contract Administration Services Management Area-San Francisco San Bruno, CA 94066	15. SECURITY CLASS. (of this report) Unclassified
	15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Unlimited

DIST[illegible] A
Appro[illegible]
[illegible]

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

Unlimited

18. SUPPLEMENTARY NOTES

Reproduction in whole or in part is permitted for any purpose of the United States Government.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Automated Decision Aids
Tactical Command and Control Decisions
Anti-Air Warfare

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report describes research conducted by Applied Decision Analysis, Inc. for the Office of Naval Research during the period from September 1980 to March 1981. The purpose of this research is to develop and implement specific automated decision aids that address current anti-air warfare (AAW) decision-making and information processing needs. The experimental aids described in this report have the capability to process and display uncertainty, and to aggregate the high volume of data characteristic of AAW command and control -

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

393952

situations into a few simplified displays. They illustrate the promise and problems associated with AAW decision-aiding systems that attempt to process uncertain subjective data in a rapidly developing command and control environment.

"Original contains color plates: All DTIC reproductions will be in black and white"

TABLE OF CONTENTS

TABLE OF CONTENTS (continued)

FIGURES

FIGURES (continued)

FIGURES (continued)

EXECUTIVE SUMMARY

This report describes some new designs for AAW decision aids. The aids discussed here are experimental systems designed to complement existing decision aids like the Naval Tactical Data System (NTDS). The proposed aids summarize and display uncertainties inherent in AAW information, and aggregate the high volume of data characteristic of AAW command and control situations into a few simplified displays. They produce displays of the probability that attacking aircraft or missiles will penetrate a task force's (battle group's) defenses undetected, the likely attack routes, the probability that a threat is located or will first appear at any location, and the expected damage levels from detected and undetected threats. Since the experimental decision aids have been formulated and partially implemented in a six month period, they are not fully developed. However, they indicate the promise and problems associated with AAW decision-aiding systems that attempt to process uncertain, subjective data in a rapidly developing command and control situation.

The AAW decision environment places a number of stringent conditions on the nature and performance of decision aids. A large number of related decisions must be made in several locations during a short period of time, and these decisions and their consequences must be communicated rapidly among numerous decision makers. There is insufficient time to formulate a unique analysis of each decision problem, and even the amount of information transferred back and forth between an aid and its users must be limited to those items that can be entered, processed, and interpreted very quickly. With these constraints, AAW decision aids cannot have the flexibility to address directly every choice that AAW personnel may be required to make. Flexibility is maintained by having AAW aids summarize the implications of available data and assessments, and allowing AAW personnel to integrate any unique elements of the situation into their decisions.

The Analytic and Information Processing Needs for AAW Decisions

In previous research for ONR, ADA has examined the AAW decision tasks performed at the unit (i.e., single ship) and task force (battle group) levels. *The decision aids described here are designed* to support AAW personnel at both the unit and task force levels, although most of them are oriented toward higher-level decisions that require the user to aggregate information about a large number of threats.

As part of this earlier research, ADA developed a taxonomy of the information processing and decision analytic functions relevant to AAW decisions. These functions include: acquiring and interpreting information; restructuring and summarizing data; identifying and analyzing the the major elements of a decision; identifying patterns and relationships in data; determining the implications of an analysis; and integrating the results of analysis with existing knowledge and intuition. The relative importance of these functions for each AAW decision task was assessed and compared to the ability of AAW personnel to perform them with existing decision aids. Those decision-making functions that are both important for one or more AAW decision tasks and difficult for AAW personnel to do with existing aids were identified as functional requirements for new decision aiding systems.

Of the functional requirements identified in this manner, the experimental aids developed by ADA address the following:

- assess and communicate the uncertainty and credibility of information produced by detection, tracking, assessment, and priority setting activities;

- combine the information available about threats and defensive weapons into an aggregate estimate of capabilities and intentions;

- sort and categorize threats according to criteria appropriate to a specific combat situation; and
- evaluate alternative defensive actions.

The research described in this report focused on AAW decision aids with two major capabilities: processing and displaying uncertainty, and aggregating track-specific data. These capabilities are central to the functional requirements for AAW decision aids listed above, and they address two of the most difficult aspects of making AAW decisions: coping with high volumes of AAW data, and assessing the uncertain implications of that data. At present, both of these tasks are done manually by AAW personnel who interpret the engagement-status information displayed by the Naval Tactical Data System (NTDS).

Processing Uncertainty

Much of the information processed by AAW decision aids is uncertain. The tracking and identification tasks are characterized by uncertain and incomplete data, but NTDS provides only a limited capability for communicating this uncertainty through an automated display to those who must make resource allocations. Inaccurate data and subjective interpretations of incomplete data are rapidly disseminated to NTDS users without a means for them to judge the confidence they should place in the resulting displays.

Another type of uncertainty relevant to AAW decisions is the implication of current tactical data for future events and decisions. AAW personnel routinely estimate the future course and tactics of tracks from current tactical data and assessments of the enemy's targets and information state (i.e., whether the enemy knows the location of various targets). Similarly, the chance that a sensor will fail to detect a threat or a defensive weapon will fail to destroy it should be considered when taking defensive actions.

In order to make these estimates and inferences, AAW personnel must mentally process several interrelated uncertainties.

The experimental decision aids developed during this research project attempt to help users make such probabilistic inferences. They deal directly with the uncertainties inherent in AAW engagements, ranging from enemy capabilities and intentions to the consequences of alternative allocations of defensive resources. The logic and displays used in these aids provide a method of processing and communicating subjective assessments of uncertain parameters and their implications for future actions.

Aggregating Data in AAW Decision Aids

Most of the data displayed by NTDS is tied to individual tracks. This means that the amount of information about an AAW engagement increases as the number of tracks increases. As more and more tracks are processed in a high-density engagement, AAW personnel will reach a point where they cannot deal adequately with all of the track-specific data for which they are responsible. When this happens, they must selectively ignore some of the data and focus their attention on only the most important pieces of information. As the amount of track-specific information increases, so does the difficulty of assessing and focusing on the important elements of the situation.

Including measures of uncertainty in AAW data and their implications for future actions and events effectively increases the amount of information stored in AAW data bases. To give AAW personnel access to this information, methods are needed to overcome the information transfer limit between the displays and the users. Two approaches that are currently in use are selective data retrieval and multiple displays. This report describes another approach based on displays that are not track specific.

The primary alternative to track-specific data in real-time AAW displays are symbols representing information about groups of tracks. NTDS uses this approach to represent a formation of aircraft with a single symbol. Further aggregation of AAW data can be accomplished by using symbols to represent such quantities as: areas where defenses are relatively weak, high concentrations of threats, probable attack routes, levels of expected damage, etc. This type of information summarizes the status of an engagement, and helps a user prioritize the processing of track-specific data.

When information about groups of tracks is combined, the result can be displayed either as a single symbol or as a quantity defined over the entire area covered by the display. The second approach is appropriate for information derived from all of the tracks in the area covered by the display, such as the probability that a threat can penetrate to any location near the task force. Quantities defined over a region and not specific to a single location can be represented by contours, colors, or shading over an entire display. Since this type of display is not used currently, it could be superimposed on the track-specific symbols generated by NTDS.

In general, the decision aids described in this report display information in the form of quantities defined over the entire area around the task force. The form of the display is a set of contours to describe the level of that quantity in each region. The interior of each contour may be shown in color to make it easier to interpret the display.

Displays of functions defined over an area were selected because they show promise for aggregating a lot of information in a form that is fairly easy to interpret, and because they are easily extended to represent the uncertainty associated with events that that can occur at any location (e.g., the probability that an undetected threat will first appear at any point near the task

force). In addition, a decision aid based on contours can *summarize* a situation with approximately the same amount of information regardless of the number of tracks present, thus circumventing the information transfer limit between displays and users.

Overview of the Proposed Aids

The experimental aids are designed to answer several related questions relevant to AAW decisions. Answers to all of these questions depend on the current defensive posture and conditions affecting sensor performance (including environmental conditions, electronic countermeasures, jamming, etc.).

- What is the probability that a threat in the vicinity of the task force will be detected by a given configuration of sensors under current or anticipated conditions?
- Which attack routes are enemy aircraft and missiles likely to use in attempting to penetrate our defenses *undetected? Which would give* them the best chance of avoiding detection?
- What is the maximum probability that a threat could penetrate undetected to any point near the task force, given estimates of the enemy's objectives and his state of information about our defenses?
- What is the probability that a currently undetected threat is located at any point near the task force, given our current defensive posture and any prior information about the threat's location (e.g., data describing a lost track)?
- What is the probability that a threat will be detected for the first time at any point near the task force, given our current defensive posture and state of information?
- Given a policy for allocating defensive weapons to threats, what is the probability that a threat can penetrate to any point near the task force without being destroyed? (The answer to this question depends on both our ability to detect threats and our ability to intercept and destroy them.)

- Given a weapons allocation policy, what is the expected damage to ships in the task force from both detected and undetected threats, and which threats are expected to cause the damage?

Aids have been implemented that would help AAW personnel answer the first five questions. Aids relevant to the last two questions have been formulated and the formats of their displays have been specified.

Aids that address these questions can be used to support many of the AAW decision-making activities at the unit and task force levels. However, these aids are best suited to activities that require integrating information about numerous tracks or anticipated enemy actions in order to establish an appropriate defensive posture. Thus, at the task force level they are likely to be most useful for the force positioning and resource allocation activities, including the positioning of ships and aircraft to maximize the chances of detecting and intercepting air threats and the assignment of defensive aircraft (interceptors and electronic surveillance) to sectors. At the unit level, they would be of the most value to AAW personnel involved in establishing priorities for dealing with threats (i.e., which threats should be engaged first).

In the remainder of this summary, each of the aids developed during this project is described briefly. Several of the aids have been implemented on a small computer graphics system; initial designs have been specified for others, but further research will be needed to implement them. All of the aids, including the ones that have been implemented, are experimental and variations of each display and alternative algorithms to generate them have been explored. For many of the proposed aids, examples of displays (the results of algorithmic testing on small sample problems) are presented in the main report. The data used in all of these sample displays is hypothetical, and not intended to represent an actual combat environment.

Detection Probabilities

This decision aid displays a measure of the probability that a combination of sensors (primarily radars) can detect a particular type of threat at any location near a task force. The boundaries of colored regions in the display represent contours of equal detection rates. On an intuitive level, regions with high detection rates correspond to areas with better radar coverage. If the detection rate in a region is doubled, then the probability that a threat can traverse the region undetected is halved. This type of display should be most useful for task force AAW decision-makers concerned with positioning defensive resources and establishing electronic emissions control strategies.

The detection rates are a function of the characteristics of the threat, the physical environment, electronic countermeasures and jamming used by both sides, and the status of various surveillance systems. The logic required to calculate the detection rates from physical data has not been developed as part of this project, and a separate research effort will be needed to specify appropriate algorithms. These algorithms should be based on empirical data describing the performance of each surveillance system under various operating conditions.

Optimum and Likely Attack Routes

This aid produces two related displays: one showing the optimum routes for attacking aircraft or missiles trying to reach any point near the task force undetected, and a second showing regions containing likely, but not necessarily optimal, attack routes. These displays are generated from the detection rates described above, and from intelligence estimates of the enemy's objectives and knowledge of the task force's detection capabilities. The displays also have the capability to reflect intelligence estimates of the direction from which the anticipated attack will be launched.

It would be unwise to concentrate too *many defenses on the* optimal attack routes. If the enemy's objectives and knowledge differ from our estimates, or the enemy decides to take a suboptimal route, defensive capability may be needed elsewhere. Therefore, the decision aid produces a display showing the relative effectiveness of various paths to a given target. Under the assumption that an enemy is more likely to take routes that minimize the probability of detection or the distance traveled, this display indicates the relative likelihood that the enemy will select routes that pass through each region. The interpretation of *this display is* consistent with the mixed (i.e., probabilistic) strategies that would result from a game theoretic approach to the problem.

Surveillance Penetration Probabilities

This decision aid produces a display of the maximum probability that a threat can penetrate undetected to any point in the vicinity of the task force. In addition to the surveillance penetration probabilities for threats that have never been detected (future attacks), the decision aid can display them for a lost track. In either case, the display is based on a "worst case" analysis using the optimum attack routes determined by the preceding aid.

Both the display of detection probabilities and this aid show the implications of positioning surveillance systems and establishing an emissions control policy. However, there is a major difference between the two displays. The latter also shows the implications of our assessments of enemy tactics, objectives, and knowledge of our current defensive posture.

Location Probabilities

This decision aid is closely related to the previous one, but it produces a display of current, rather than projected, threat status.

The display shows the probability that one or more undetected threats are currently located at any point in the vicinity of the task force. The display can be generated for threats that have never been detected and for lost tracks, but different assumptions and data are needed for these two situations.

Once the location probabilities have been calculated for several lost tracks (or tracks that have never been detected), it is relatively simple to calculate and display the probability that one or more undetected threats exist at any location. Displays for the location probabilities associated with multiple lost tracks have not been implemented, but they are a direct extension of the single-track displays implemented during this research project.

First Detection Probabilities

This decision aid displays the probability that a currently undetected track (usually part of an anticipated future attack) will be detected for the first time at any location near the task force. This information is closely related to the surveillance penetration and location probabilities. However, first detection probabilities are better suited to the needs of those assigning defensive weapons to threats because they indicate the areas where defensive weapons may have to be targeted. First detection probabilities also determine the reaction times within which AAW personnel will have to assign and fire defensive weapons. For this reason, first detection probabilities are used by other decision aids to calculate the expected consequences of an engagement.

First detection probabilities can also be calculated and displayed for a lost track. However, this requires considerably more computation than the case of a track that has never been detected; as a result, this algorithm has not been implemented. Alternative, more efficient algorithms to deal with this case are under investigation.

Total Penetration Probabilities

This aid displays the probability that a threat can penetrate to any point near the task force without being destroyed. In order to calculate this quantity, the aid combines the first detection probabilities produced by the preceding aid with information about known (i.e., detected) threats and a policy for allocating defensive weapons to threats. A simple combat model describes the effectiveness of defensive weapons. The aid produces a display similar to that for surveillance penetration probabilities, but shows the probability of survival rather than nondetection. The inputs, displays, and some of the logic for this aid have been specified, but it has not been implemented.

This decision aid does not attempt to optimize the allocation of defensive weapons. Instead it shows the effect of a defensive policy established by the user. For instance, the user might specify the defensive weapons to be used against threats in various areas near the task force, and require that a single defensive weapon be allocated to each threat within a given distance of the task force, followed by a second weapon if the first fails. The decision aid would show the implications of this policy for threat penetration; if the results are unacceptable, the user could try a different weapons allocation strategy. In this manner, AAW personnel would retain control over all of their resource allocation decisions, and the decision aid would serve to help them test alternative policies and select an acceptable one.

Sources of Expected Damage

This aid produces a display of the expected damage (i.e., threat level) associated with both known tracks and undetected threats. The display is generated in real time, and it changes as threats are detected, move to new positions, or encounter defensive weapons. It

is designed to support AAW personnel at the unit level who prioritize threats and assign defensive weapons to counter them.

Like the preceding display, this one has not been implemented, but its inputs and the form of the display have been specified. It includes both track-specific information (for detected threats) and a threat level defined over the entire area (for undetected threats). Associated with each known (i.e., detected) threat is the expected amount of damage it can cause ships in the task force. In the display the expected damage associated with detected threats is represented by colored symbols for the tracks. For undetected threats, the expected damage is represented by contour levels or colored regions around the task force.

Directions for Future Research

A significant amount of work remains to be done in order to complete development of the proposed aids. In particular, implementation of the displays of location probabilities for undetected (anticipated) tracks and first detection probabilities for lost tracks has been deferred because these displays require more computation time than other displays. Alternative algorithms are currently under investigation, including an approach based on Markov processes described in the Appendix.

Two of the decision aids discussed in this report have not been implemented: the displays of total penetration probability and sources of expected damage. However, the current research project has resulted in a specification of the required input data and the form of each display. One goal of future research will be to complete the design and implementation of these displays. A major requirement for production of these two displays is the development and implementation of a simple combat model.

Continuing development of the aids presented in this report should include the implementation of sensor detection models based on empirical data. In addition, the experimental aids should be tested on a larger computer system, so that algorithmic inefficiencies can be assessed and corrected.

Displays showing the density of detected, lost, and undetected (anticipated) threats, as well as the density of friendly forces, represent another possible avenue of future research. Such displays could lead in a natural way to status summaries of controlled areas and battle zones. This type of display could be valuable in showing where each side is vulnerable to attack or the repositioning of opposing forces.

I. INTRODUCTION

This report describes some new designs for AAW decision aids. The aids discussed here are experimental systems designed to complement existing decision aids like the Naval Tactical Data System (NTDS). The proposed aids summarize and display uncertainties inherent in AAW information, and aggregate the high volume of data characteristic of AAW command and control situations into a few simplified displays. They produce displays of the probability that attacking aircraft or missiles will penetrate a task force's (battle group's) defenses undetected, the likely attack routes, the probability that a threat is located or will first appear at any location, and the expected damage levels from detected and undetected threats Since the experimental decision aids have been formulated and partially implemented in a six month period, they are not fully developed. However, they indicate the promise and problems associated with AAW decision-aiding systems that attempt to process uncertain, subjective data in a rapidly developing command and control situation.

The AAW decision environment places a number of stringent conditions on the nature and performance of decision aids A large number of related decisions must be made in several locations during a short period of time, and these decisions and their consequences must be communicated rapidly among numerous decision-makers. There is insufficient time to formulate a unique analysis of each decision problem, and even the amount of information transferred back and forth between an aid and its users must be limited to those items that can be entered, processed, and interpreted very quickly. In this environment AAW aids must help a user focus on the most significant elements of the problem and reach his own conclusions about the best course of action.

Time constraints on AAW decisions require that much of the logic and format of the decision-making process be specified before any decisions take place. This means that AAW decisions aids must be based on prepackaged algorithms and display modes, with only a minimum of subjective input from users. The number of assessments must be limited to a few key parameters that do not change rapidly throughout the course of an engagement. With these constraints, AAW decision aids cannot have the flexibility to address directly every choice that AAW personnel may be required to make. Flexibility is maintained by havin~ AAW aids summarize the implications of available data and assessments, and allowing AAW personnel to integrate any unique elements of the situation into their decisions.

Organization of This Report

The remainder of this introduction summarizes ADA's previous investigation of the functions that should be provided by AAW decision aids, and the approach taken in this research project to design new aids that provide some of these functions. Section II discusses the benefits and difficulties associated with processing uncertainty and aggregating data associated with AAW decisions. This section concludes with a discussion of alternative types of displays for decision aids that accomplish these tasks, and describes the type of displays selected for aids developed during this research project. Section III describes each of the decision aids in detail, shows how they are related, and discusses their data requirements, assumptions, displays, and limits. Section IV describes areas where further research is needed to overcome some of the limitations associated with the experimental decision aids, and to develop improved aids capable of supplying additional support for AAW decision making activities. The Appendix specifies the logic and mathematics contained in each algorithm used by the aids. In some cases alternative algorithms are described, together with their limitations and relative advantages.

AAW Decision-Making Activities

In previous research for ONR, ADA has examined the AAW decision tasks performed at the unit (i.e., single ship) and task force (battle group) levels. The decision aids described here are designed to support AAW personnel at both the unit and task force levels, although most of them are oriented toward higher-level decisions that require the user to aggregate information about a large number of threats

The decision-making activities at the task force level deal primarily with assigning AAW responsibilities and assets to individual units, and then coordinating their activities. The major decision-making activities that occur at the task force level are:

- establish operational guidelines for AAW personnel and interpret rules of engagement specified by higher levels of command (e.g., define the conditions under which a unit can fire on an air target without further authorization);
- interpret intelligence reports and anticipate the type of air attack that might occur;
- position ships and aircraft to maximize the chances of detecting and intercepting air threats;
- establish controls on electronic emissions (EMCON' and select electronic warfare tactics;
- designate air defense sectors and assign responsibility for AAW operations within a sector to a ship or a sector AAW coordinator;
- assign defensive aircraft (interceptors and electronic surveillance) to sectors;
- establish procedures for the movement of friendly aircraft near the task force; and
- resolve conflicting decisions or actions by AAW units and coordinate their activities.

The last decision-making activity in this list is made both before and during an engagement by one of several officers designated to coordinate air defense activities. The principal officer in this category is the antiair warfare coordinator (AAWC). He is supported by a force track coordinator (FTC) who monitors and reconciles the tracking and identification decisions made by units in each sector. Since the decisions made by these officers are closely related to those made at the unit level, they currently are supported by the same decision aids. However, these aids provide little direct support for the other AAW decision activities in the list above.

The decision-making activities that occur at the unit level are:

- detect and track aircraft and missiles;
- identify each track (i.e., determine the type of aircraft or missile, and whether it is friendly or hostile);
- assess the degree of danger posed by a threat (i.e., estimate its mission and the likelihood it will succeed);
- establish priorities for dealing with threats (i.e., which threats should be engaged first);
- assess the capabilities of alternative weapon systems for countering a threat (i.e., determine whether a weapon can intercept the threat and the likelihood that it will stop the threat);
- assign defensive weapons to counter each threat; and
- decide when, or under what conditions, to fire defensive weapons (e.g., fire a missile when the threat reaches a certain position).

In practice, these activities are treated as separate decisions, even though many of them are information processing activities that support the final decisions to assign a defensive weapon to an air target and fire the weapon.

At the unit level, all of the decision-making activities take place in a combat information center (CIC), where they are supported by either automated or manual decision aids. The central decision maker in the CIC is the tactical action officer (TAO). The TAO is trained to manage all unit AAW operations and coordinate his actions directly with the AAWC.

The remaining AAW personnel in a ship's CIC report to the TAO. Their responsibilities correspond to the unit level decision activities listed above. On a ship with an automated AAW system, detector/trackers watch radar screens and assign synthetic (i.e., computer generated) symbols to radar signals reflected by aircraft or missiles. As the radar signals move across the screen, the detector/trackers instruct the computer to move the synthetic symbols along the same paths. ID operators use a variety of information to estimate the identity of each object on the screen, and modify the synthetic symbols so others can see what they represent.

Assessing and prioritizing the threats represented by the NTDS symbols is primarily the job of the TAO and, if the ship has an automated AAW system, the ship's weapons coordinator (SWC). The SWC keeps track of the availability of defensive weapons systems and helps the TAO assess their capability to destroy air threats. Defensive weapons can be assigned to threats by either the SWC or the TAO, with the TAO exercising veto power if he does not accept the SWC's decisions. Working with the SWC may be several specialized weapon controllers, including an intercept controller, fire control system controller, and engagement controller, depending on the size of the ship and the number of stations in the CIC. These operators implement the decisions of the SWC and TAO using the information entered by trackers and ID operators.

The primary automated decision aid for AAW personnel is the Naval Tactical Data System (NTDS). NTDS is a specialized information processing system that accepts real-time data from radars, aircraft, weapons systems, operators, and other ships; processes this information and displays it on small-screen consoles; exchanges the information via digital data links with other NTDS units (or units with compatible data systems); and reports some of the information to other commands and non-NTDS units via teletype. The primary control and output device for this sytem is an NTDS console, which consists of a small screen resembling a radar repeater and a variety of control keys. There are several types of NTDS consoles, and most consoles can be operated in several different modes. However, the consoles use common data definitions and symbology, and information or instructions entered at one console can be observed at other consoles. For instance, a track detected using one ship's radar can be entered into an NTDS console and it will appear immediately on consoles aboard other ships with NTDS systems that are connected to the same digital data net.

The Analytic and Information Processing Functions Needed to Support AAW Decisions

In previous research for ONR, ADA developed a taxonomy of the information processing and decision analytic functions relevant to AAW decisions. These functions include: acquiring and interpreting information; restructuring and summarizing data; identifying and analyzing the the major elements of a decision; identifying patterns and relationships in data; determining the implications of an analysis; and integrating the results of analysis with existing knowledge and intuition. The relative importance of these functions for each of the AAW decision tasks listed above was assessed and compared to the ability of AAW personnel to perform them with existing decision aids. While these assessments were not precise, they indicated general areas where appropriate decision aids are

most likely to improve the decision process. Those decision-making functions that are both important for one or more AAW decision tasks and difficult for AAW personnel to do with existing aids were identified as functional requirements for new decision aiding systems.

The functional requirements identified in this manner are:

- assess and communicate the uncertainty and credibility of information produced by the detection, tracking, assessment, and priority setting activities;
- combine the information available about threats and defensive weapons into an aggregate estimate of capabilities and intentions;
- store and recall subjective assessments of the status and capabilities of threats and defensive weapons;
- sort and categorize threats according to criteria appropriate to a specific combat situation;
- identify predefined patterns of data that indicate the existence and identity of tracks, and the capabilities of threats;
- identify trends in enemy tactics; and
- predict outcomes and evaluate alternative defensive actions.

The experimental aids described in this report deal directly with the uncertainties inherent in AAW engagements ranging from enemy capabilities and intentions to the consequences of alternative allocations of defensive resources. The logic and displays used in these aids provide a method of processing and communicating subjective assessments of uncertain parameters and their implications for future actions.

Even though AAW personnel have a limited amount of time to make and communicate subjective assessments of the performance of AAW systems, it is important to design AAW decision aids that can

process and display the implications of this information as well as physical data. Currently subjective information is transmitted verbally among AAW personnel and used to interpret NTDS displays. For example, a TAO may discount some of the information displayed on an NTDS console because he knows it was entered by personnel with relatively little training or experience. Alternatively, he may take precautionary actions if he suspects that the data on an NTDS display is incomplete and threats exist that are not displayed. If future decision aids are to provide more direct support for AAW decision-making activities, they must process and draw implications from the uncertainties and subjective estimates considered by AAW personnel. The aids described in this report represent a first step in this direction.

The Approach Used to Design the Experimental Aids

The research described here started with a re-examination of the AAW decision functions that could benefit most from improved decision aids. The list of functional requirements described above was consolidated into a few basic functions that could be provided by a small set of experimental aids. For reasons discussed in the next section, the research focused on aids capable of processing uncertainty and aggregating the high volumes of data characteristic of AAW command and control. Alternative methods of summarizing uncertain data were considered, and several promising displays were selected for further development.

An initial specification of the data processing required to generate the displays indicated that it could be done with a small computer if certain assumptions could be made about the relationships among the uncertain quantities shown in the displays. These assumptions are discussed in Section III. The proposed displays were discussed with Navy research personnel involved in the design of AAW systems, and revised in accordance with their

comments. Finally, some of the experimental decision aids were implemented using a small computer graphics system. This turned out to be a very important step since some of the difficulties associated with using and interpreting the aids were not apparent until we tried to implement them. For instance, attempts to describe the user's prior state of information about an undetected track in terms of attack routes rather than track locations resulted in counterintuitive and inconsistent displays. This problem would not have been apparent without an automated display generator. By experimenting with the algorithms and displays, we were able to identify and correct problems in interpreting the inputs and outputs of the decision aids. (The significance of prior information about an undetected track is discussed in Section III.)

II. PROCESSING UNCERTAINTY AND AGGREGATING DATA IN AAW DECISION AIDS

This research project focused on designing AAW decision aids with two major capabilities: processing and displaying uncertainty, and aggregating track-specific data. These capabilities are central to the list of functional requirements for AAW decision aids described in the previous section, and they address two of the most difficult aspects of making AAW decisions: coping with high volumes of AAW data, and assessing the uncertain implications of that data. At present, both of these tasks are done manually by AAW personnel who interpret the engagement-status information displayed by NTDS.

Most of the data displayed by NTDS is tied to individual tracks. This means that the amount of information about an AAW engagement increases as the number of tracks increases. As more and more tracks are processed in a high-density engagement, AAW personnel will reach a point where they cannot deal adequately with all of the track-specific data for which they are responsible. When this happens, they must selectively ignore some of the data and focus their attention on only the most important pieces of information. However, as the amount of track-specific information increases, so does the difficulty of assessing and focusing on the important elements of the situation.

This problem will not be overcome by storing and displaying additional track-specific information. For instance, giving each track a score or symbol indicating its importance (or threat level) may simply force more data on an overburdened user. One solution, which is currently in use, is to automatically direct the user's attention to the next most important track using a scoring rule to rank tracks. An alternative, which allows the user to exercise more judgement in the selection process, is to summarize the AAW situation with a display that is not tied to individual tracks. For example,

an aid could display the density of enemy aircraft or the coverage of uncommitted defensive weapons in the vicinity of a task force. The advantage of this type of display is that it can summarize the situation with the same amount of information regardless of the number of tracks present. The question is whether aggregating track-specific information and displaying it in this manner can really help AAW personnel who must ultimately deal with individual tracks.

AAW personnel will only divert their attention from track-specific data if they have an aggregate display that presents the information they were trying to deduce from individual tracks. Whether or not an aggregate display can accomplish this goal depends on the needs of various users. Those making decisions specific to a single track (e.g., vectoring an aircraft to intercept a threat) may have little need for an aggregate display, except when they are searching for the next track to process. However, those managing defensive actions against multiple threats, both at the unit and track force levels, need to integrate and understand the major elements of the situation, set priorities and allocate resources for those making track-specific decisions, and maintain a capability to deal with new threats. The aggregate displays described in this report are oriented toward this latter group of decision-makers.

At the same time that AAW personnel are being overloaded with data, they need to know the degree of uncertainty associated with the data they receive. A good defense for an AAW situation presented on an automated display may be inappropriate if the actual situation could be significantly different. For example, if there is a significant probability that an undetected threat exists near a task force, it would be prudent to assign additional surveillance systems and defensive weapons to the areas where the threat is likely to exist. Some of the displays described in this report are designed to help a user judge the likelihood and probable location of undetected threats.

Much of the information processed by AAW decision aids is uncertain. The tracking and identification tasks are characterized by uncertain and incomplete data, but NTDS provides only a limited capability for communicating this uncertainty through an automated display to those who must make resource allocations. Translating radar signals into the computer-generated symbols displayed by NTDS requires a sequence of subjective judgements or automated inferences that cannot always be correct. Inaccurate data and incorrect interpretations of incomplete data are rapidly disseminated to NTDS users without a means for them to judge the confidence they should place in the resulting displays.

Another type of uncertainty relevant to AAW decisions is the implication of current tactical data for future events and decisions. AAW personnel routinely estimate the future course and tactics of tracks from current tactical data and assessments of the enemy's targets and information state (i.e., whether the enemy knows the location of various targets). Similarly, the chance that a sensor will fail to detect a threat or a defensive weapon will fail to destroy it should be considered when taking defensive actions. These estimates and inferences require AAW personnel to process mentally several interrelated uncertainties. The decision aids developed during this research project attempt to help users make these probabilistic inferences.

Limits on the Amount of Information Displayed by a AAW Decision Aid

The problem with displaying the uncertainties associated with current AAW data or its implications for future events is that AAW systems already have the capability of displaying more information than users can interpret in the time available. Showing AAW personnel the degree of uncertainty associated with displayed data will require them to deal with additional information. For instance, a display showing the uncertainty in a track's location might require

the user to visualize the track as a shaded region or a set of contours rather than a single symbol. This may not be practical in a situation where many tracks must be displayed at once.

The solution to this problem lies in aggregating track-specific data into displays that are independent of the number of tracks. Uncertainty and aggregation are complementary requirements; one adds to the complexity of a display while the other reduces it. However, both have the potential for transferring to decision aids some of the information processing and analytic functions currently being done by AAW personnel.

A simplified view of the problem with displaying uncertainty in AAW data is shown in Figure 1. Information about an AAW engagement, including radar signals, weapons status, intelligence reports, location of friendly forces, etc., is collected by sensors and data links and stored in a data base. In a typical AAW situation, the contents of the data base are constantly updated to reflect the physical environment. Information is extracted from the data base and displayed to AAW personnel, who use it as the basis for making defensive resource allocations. The amount of information that can be stored in the data base and the amount of processing that can be done in real time to generate the displays depends on the computer equipment available. These limits can be expected to improve as new technology becomes available. However, there is also a limit on the rate at which AAW personnel can interpret and respond to the information contained in the displays. The displays must be designed to allow personnel to operate within this limit.

Measures of uncertainty in AAW data and their implications for future actions and events effectively increase the amount of information stored in the data base. To give AAW personnel access to this information, methods are needed to overcome the information transfer limit between the displays and the users. Two approaches

FIGURE 1

FLOW OF INFORMATION IN AN AUTOMATED
AAW DECISION AID

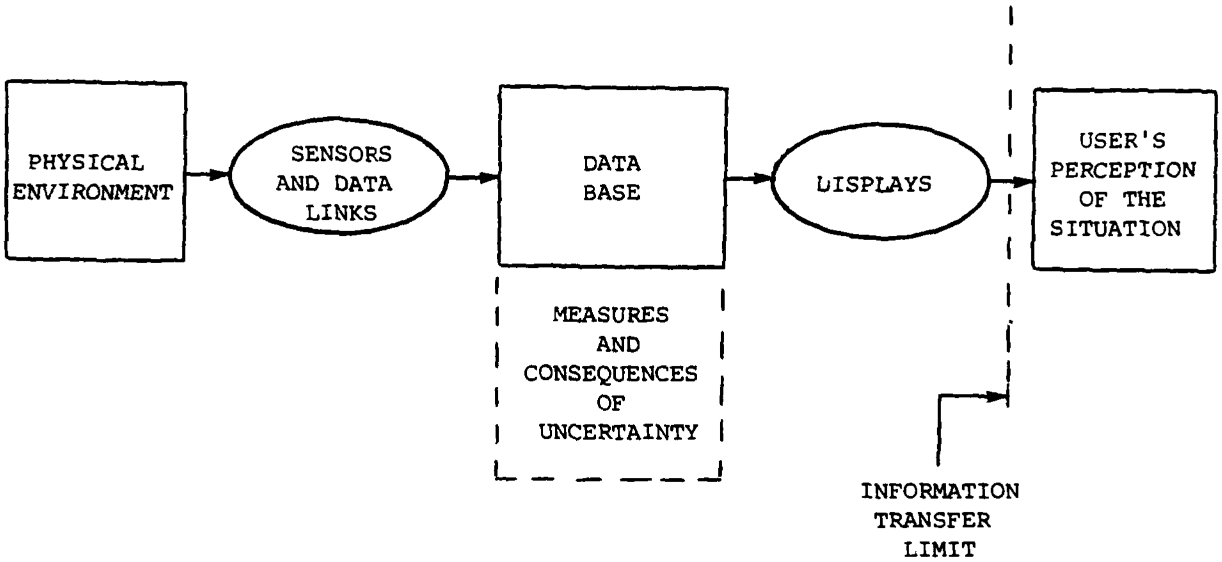

If measures of uncertainty and their consequences are computed and stored in the data base, then the displays must be designed so this information can be communicated to users without exceeding the rate at which they can interpret the information.

that are currently in use are selective data retrieval and multiple displays. For instance, an NTDS operator can limit the information on his display to symbols representing only major threats, and he can request additional information about any track to be presented on panels located near the main display. Since AAW personnel already have selective data retrieval and multiple displays, and are having difficulty dealing with all of the significant data available, giving them the option of displaying the uncertainty in track-specific data will not help them overcome the limit in the amount of information they can interpret. Displays of the uncertainty in AAW information will be useful only if the user's attention is restricted to a small number of tracks or if the track-specific data is first aggregated into a relatively simple display.

Aggregating detailed AAW data into displays that are easily interpreted means that some of the subjective information processing currently being done by AAW personnel must be performed by decision aids. Whether or not decision aids deal with uncertainty, increasingly complex decision problems will require that more of the work of interpreting AAW information be done by decision aiding systems. However, it will be difficult to accomplish this without giving the aids the ability to process uncertainty. Many of the major interpretive tasks performed by AAW personnel involve assessments of uncertain quantities (e.g., consequences of alternative defensive actions).

However, decision aids that process uncertain, subjective information to produce aggregate displays of AAW data will be more difficult to understand and verify than aids based on "objective" calculations. For instance, it is relatively easy to have confidence in an aid that extrapolates a track's trajectory. However this information may have less bearing on defensive resource allocations than the likelihood that the threat will try to reach various

targets, including those not located along that trajectory. An aid that displays these likelihoods would require a set of assumptions or subjective estimates that users would have to understand and accept.

To assume more of the work of data interpretation and aggregation currently performed by AAW personnel, aids will have to incorporate the assumptions and heuristics people use to interpret NTDS displays and reach a decision. AAW personnel will have confidence in heuristic aids only if they have control over the assumptions they contain. To allow this type of control, the major assumptions contained in the experimental aids are expressed as variables that can be modified by users to reflect their assessment of an AAW situation.

Methods of Aggregating AAW Information and Displaying Uncertainty

In order to aggregate track-specific information in AAW decision aids, it is necessary to define the manner in which the composite information should be displayed to users. The central display in current AAW systems is a computer-generated map showing the location of attacking and defensive units. The spatial orientation of this display reflects the physical nature of the problem and makes it easy to interpret the relationships among the symbols shown on the screen. While other types of displays are possible (e.g., a display showing expected attrition as a function of time), it would be desirable for any new aids to retain the spatial orientation of NTDS displays for compatibility and ease of interpretation. The displays described in Section III all present data as a function of its location relative to the task force.

One method of aggregating data that has been incorporated in NTDS is a replay of recent track activity in accelerated time. This type of display does not really integrate the track-specific information, but it helps the user do so. The same sort of display

could be used to show projected track movements using a deterministic algorithm to calculate future track positions. The advantage of this display is that it uses the visual information-processing abilities of AAW personnel to give them an overview of an engagement's immediate history or projected future. However, because a user must watch the display for a period of time in order to integrate the information presented, it is not well suited to the AAW environment with its limited decision times.

The primary alternative to track-specific data in real-time AAW displays are symbols representing information about groups of tracks. NTDS uses this approach to represent a formation of aircraft with a single symbol. Further aggregation of AAW data can be accomplished by using symbols to represent such quantities as: areas where defenses are relatively weak, high concentrations of threats, probable attack routes, levels of expected damage, etc. This type of information summarizes the status of an engagement, and helps a user prioritize the processing of track-specific data.

When information about groups of tracks is combined, the result can be displayed either as a single symbol or a quantity defined over the entire area covered by the display. The first approach is appropriate when the aggregate information refers to a group of tracks located near each other, such as a symbol indicating a high density of threats in a particular area. The second approach is better suited to information derived from all of the tracks in the area covered by the display, such as probability that a threat can penetrate to any location near the task force. Quantities defined over a region and not specific to a single location can be represented by contours, colors, or shading over an entire display. Since this type of display is not used currently, it could be superimposed on the track-specific symbols generated by NTDS.

All of the decision aids described in Section III display information in the form of quantities defined over the entire area around the task force, except for part of the display showing the sources of expected damage to ships in the task force. Displays of functions defined over an area were selected because they show promise for aggregating a lot of information in a form that is fairly easy to interpret, and because they are easily extended to represent the uncertainty associated with events that that can occur at any location (e.g., the probability that an undetected threat will first appear at any point near the task force). Another reason for concentrating on this type of display is that the computer programs needed to generate the displays could be written as a general-purpose software package and shared by the various decision aids. This commonality simplified the implementation process and made it possible to test several displays and algorithms in a short time period.

III. AAW DECISION AIDS THAT AGGREGATE UNCERTAIN INFORMATION

This section discusses several experimental decision aids for processing, aggregating, and displaying uncertain information relevant to the decisions made by AAW personnel. Each aid is described in terms of its inputs, displays, data sources, general logic, required assumptions, and data processing requirements. The detailed logic and mathematics contained in the algorithms used by the aids are discussed separately in the Appendix. Some of the experimental aids have been implemented using a small computer graphics system. Initial designs have been specified for others but further research will be needed to implement them. All of the aids, including the ones that have been implemented, are experimental, and variations of each display have been explored. The variations in each aid's displays are discussed in this section, and alternative algorithms for producing the displays are discussed in the Appendix.

The aids are designed to answer several related questions relevant to AAW decisions.

- What is the probability that a threat in the vicinity of the task force will be detected by a given configuration of sensors under current or anticipated conditions?

- Which attack routes are enemy aircraft and missiles likely to use in attempting to penetrate our defenses undetected? Which would give them the best chance of avoiding detection?

- What is the maximum probability that a threat could penetrate undetected to any point near the task force, given estimates of the enemy's objectives and his state of information about our defenses?

- What is the probability that a currently undetected threat is located at any point near the task force, given our current defensive posture and any prior information about the threat's location (e.g., data describing a lost track)?

- What is the probability that a threat will be detected for the first time at any point near the task force, given our current defensive posture and state of information?

- Given a policy for allocating defensive weapons to threats, what is the probability that a threat can penetrate to any point near the task force without being destroyed? (The answer to this question depends on both our ability to detect threats and our ability to intercept and destroy them.)

- Given a weapons allocation policy, what is the expected damage to ships in the task force from both detected and undetected threats, and which threats are expected to cause the damage?

Aids have been implemented that would help AAW personnel answer the first five questions. Aids relevant to the last two questions have been formulated and the formats of their displays have been specified. Although they have not been implemented, they are summarized in this section.

Aids that address these questions can be used to support all of the AAW decision-making activities at the unit and task force levels listed in the Introduction. However, these aids are best suited to activities that require integrating information about numerous tracks or anticipated enemy actions in order to establish an appropriate defensive posture. Thus, at the task force level they are likely to be most useful for the force positioning and resource allocation activities, including the positioning of ships and aircraft to maximize the chances of detecting and intercepting air threats and the assignment of defensive aircraft (interceptors and electronic surveillance) to sectors. At the unit level, they would be of the most value to AAW personnel involved in establishing priorities for dealing with threats (i.e., which threats should be engaged first).

It is difficult to predict which of these aids will prove most useful for each of the decision-making activities listed in the Introduction, since both the questions addressed by the aids and the

activities themselves are interrelated. Individual users may prefer to look at different displays when performing the same activities, and should be given the option of doing so. The experimental aids that appear best suited to each AAW decision-making activity are shown in Figure 2. As this figure indicates, the aids dealing with the likelihood that a threat will penetrate both surveillance coverage and defensive weapons, and the resulting levels of expected damage, are relatively important for unit level activities. The other aids are better suited to positioning and resource allocation decisions at the task force level.

Similar data and logic are used by the aids associated with each question listed above. In fact, some of the aids use information calculated and displayed by others. This means that assumptions and assessments made for one aid will be reflected in the displays generated by others. It also means that the aids should be implemented in a manner that allows them to share data, algorithms, and the software that generate the displays.

The flow of logic and data used by the experimental aids is shown in Figure 3. The first aid uses empirical information about the performance of sensors under various operating conditions and real-time data about environmental conditions, including enemy jamming and ECM activities, to calculate and display a measure of the probability that threats will be detected at any location near the task force. The second aid uses this information, and estimates of the enemy's objectives and knowledge of the task force's defences, to calculate the optimum and likely attack routes. This aid displays the likelihood of alternative attack routes as a function of how closely they correspond to the estimated objectives of the enemy.

The information about detection rates and likely attack routes is combined with estimates of the size and direction of an anticipated attack, and with information about lost tracks, to

FIGURE 2

THE AAW DECISION-MAKING ACTIVITIES SUPPORTED BY THE EXPERIMENTAL AIDS

	Detection Probabilities	Attack Routes	Surveillance Penetration	Location Probabilities	First Detection Probabilities	Total Penetration	Source of Damage
Unit Level Activities							
Detect and track threats	X			X	X		
Identify tracks		X		X			
Assess threat intentions and capabilities						X	X
Prioritize threats							X
Assess defensive weapons and assign to threats				X	X	X	X
Decide when to fire						X	X
Task Force Level Activities							
Establish guidelines for engagement					X	X	
Interpret intelligence reports and anticipate type of attack		X	X				
Position ships and aircraft	X	X	X	X	X	X	X
Establish controls on electronic emissions	X		X			X	
Designate air defense sectors	X	X	X		X	X	
Assign aircraft to sectors				X	X	X	X
Establish procedures for friendly aircraft		X					
Resolve conflicting unit-level AAW decisions				X			X

FIGURE 3

THE SEQUENCE OF LOGIC AND DATA USED BY THE EXPERIMENTAL DECISION AIDS

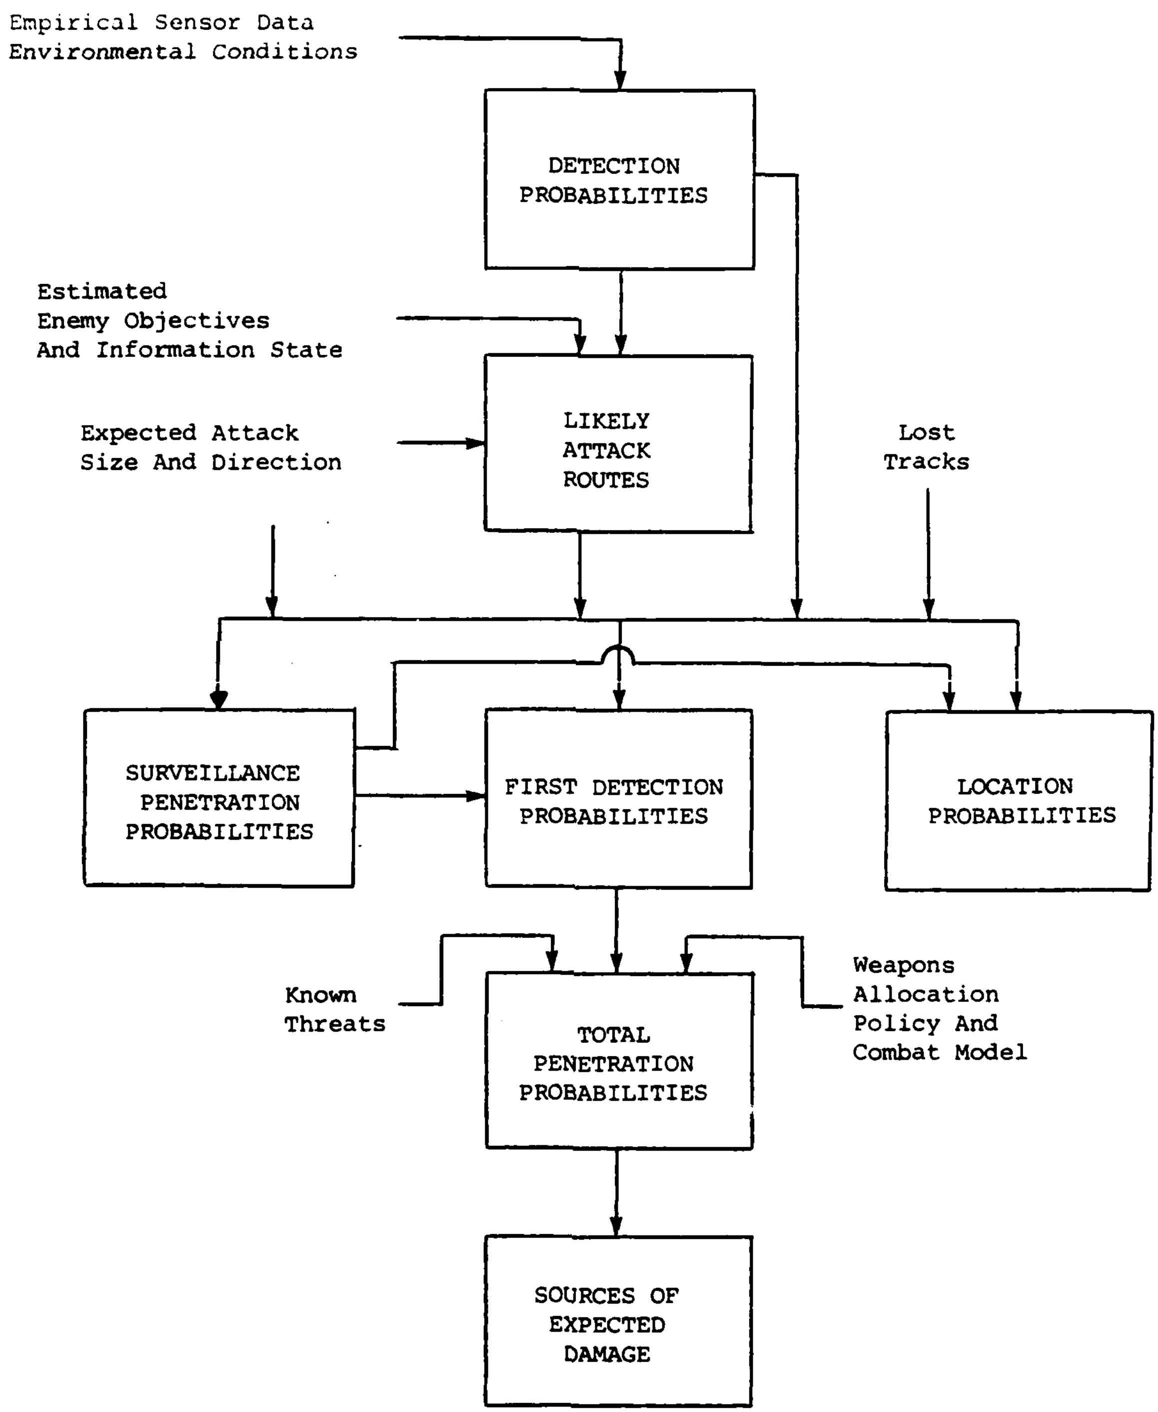

produce three related displays. These show the probability that a threat can reach any point undetected, the probability that an undetected threat is currently located at any point, and the probability that a threat will first appear at any point. The first of these displays is based on the best routes for reaching each point in the vicinity of the task force. The other two displays are based on the estimated targets of attacking aircraft or missiles. All of the displays can be generated for threats that were detected previously and then lost.

The last two decision aids deal with combat as well as detection. They display the probability that a threat can penetrate to any point in spite of the use of defensive weapons. In addition, they show the sources of expected damage to ships in the task force. These displays make use of the first-detection probabilities produced by the other aids because the effectiveness of defensive weapons depends on where a threat is first detected. These two aids process both known threats and anticipated (but currently undetected) threats, using a simple combat model and a specified policy for allocating defensive weapons. They show the effectiveness of alternative weapons allocation policies and indicate the relative hazard posed by detected and undetected threats. This information can be used to determine the defensive resources that should be held in reserve to deal with new threats as they appear.

The remainder of this section discusses each of the experimental decision aids in more detail, showing typical outputs as they appear on a small-screen computer-driven display, or as they will appear when all of the aids are implemented.

Detection Probabilities

This decision aid displays a measure of the probability that a combination of sensors (primarily radars) can detect a particular type of threat at any location near a task force. A typical display

is shown in Figure 4 for a group of three ships with active radars. The ships are represented by small circles, and the colored regions around them indicate the likelihood that a threat will be detected in a fixed time interval. For all displays presented in this report, red represents the region of highest probability, yellow areas are associated with moderately high probabilities, green zones denote moderately low probability regions, and the regions of lowest probability are colored blue. Thus, the colored regions in Figure 4 indicate that the probability of avoiding detection decreases as the threat moves closer to one or more of the ships. (If this aid were implemented on a more sophisticated computer graphics system, the display would also indicate the actual range of probabilities associated with each color.)

An alternative way to display the detection probabilties in Figure 4 is with a set of contours like those shown in Figure 5. In fact, all of the displays of quantities defined over the entire area near the task force can be represented by contours similar to those in Figure 5. Contour displays are easier to generate that colored regions, and they do not require a color graphics system. However it is more difficult for a user to interpret a contour display, especially when many additional symbols are superimposed on them. Since ease of interpretation is important in AAW displays, we have chosen to implement them with colored regions rather than contours.

The boundaries of the colored regions in Figure 4 and the contours in Figure 5 represent contours of equal detection rates. The probability that a threat will be detected at a particular point is equal to the detection rate at that point multiplied by the length of time that the threat remains at (or near) that location. As a threat moves along any given path, the probability that it will remain undetected can be calculated from the detection rate at points

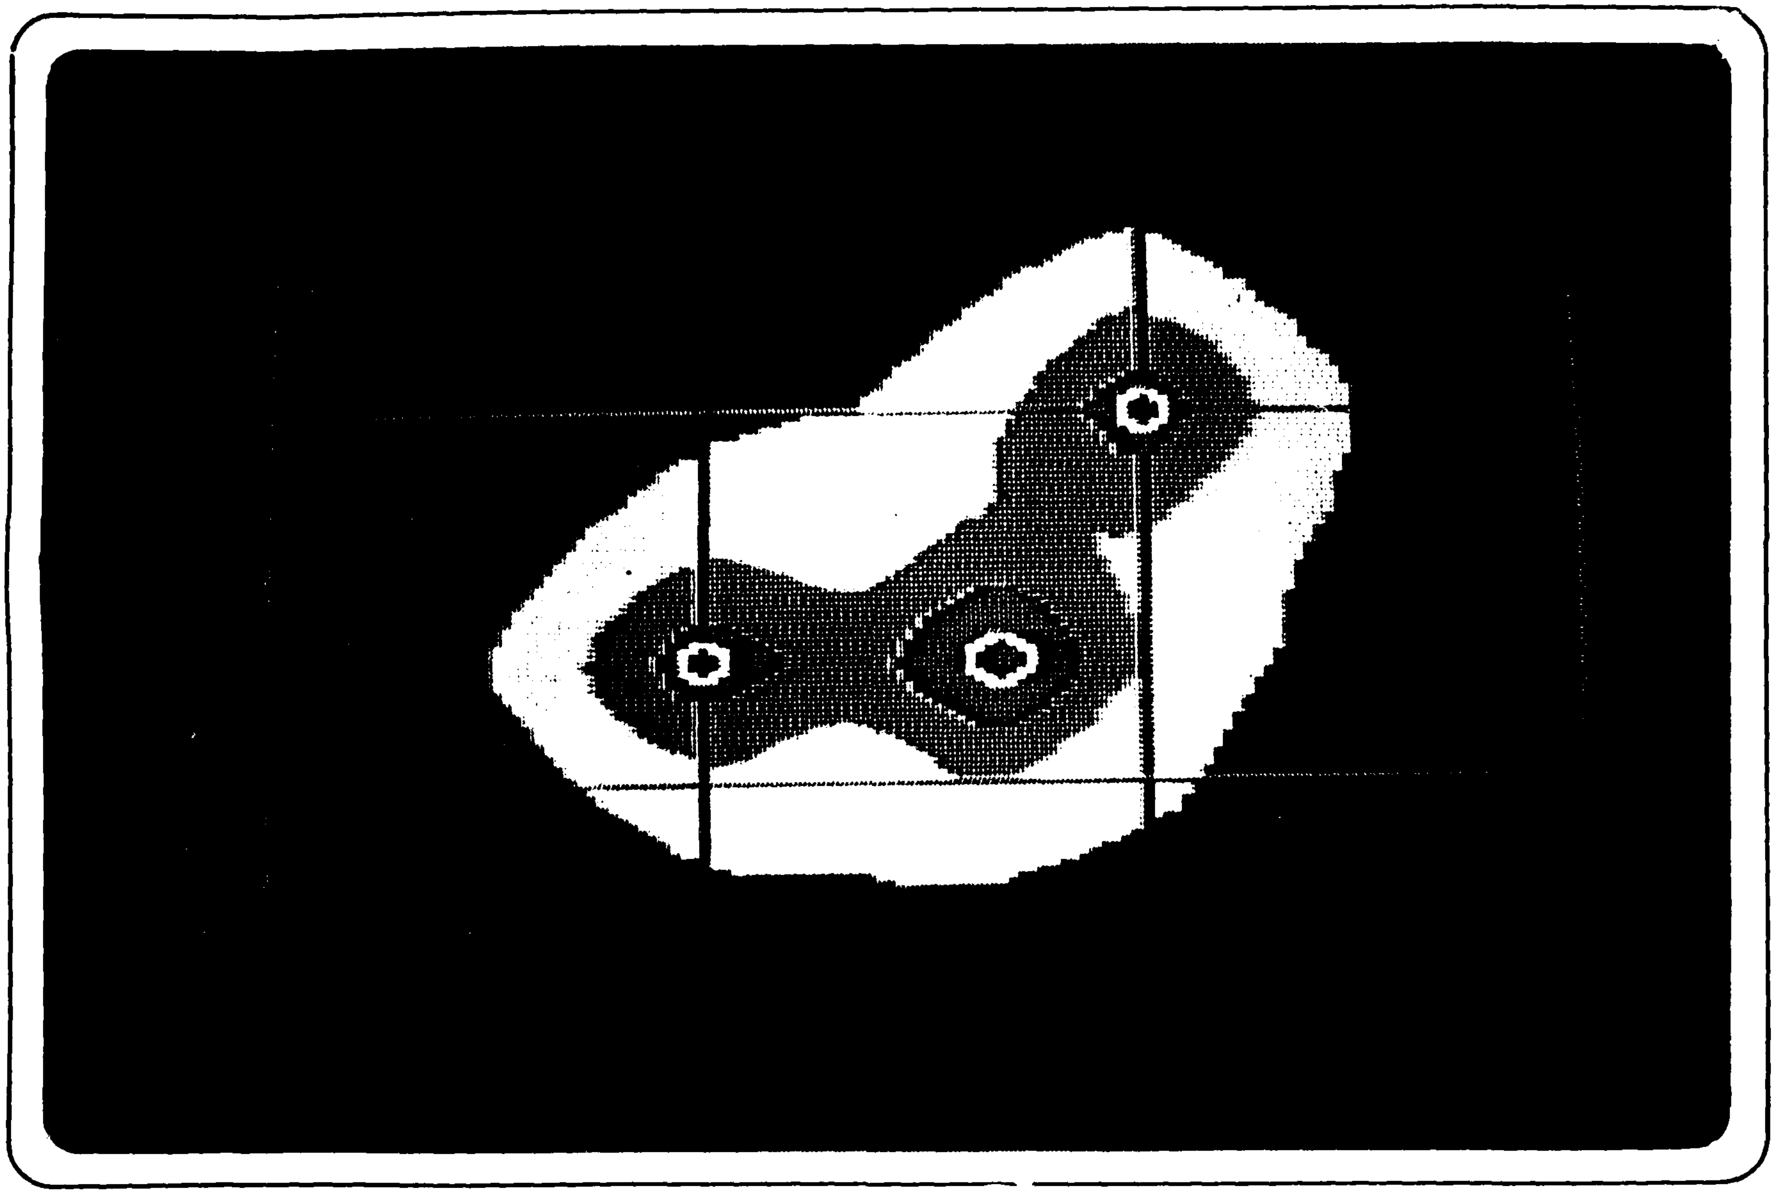

FIGURE 4

TYPICAL PROBABILITIES OF AVOIDING DETECTION FOR A GROUP OF THREE SHIPS

The distance between the grid lines i[illegible] ly 150 miles. Red corresponds to regions of highest probability, and blue corresp[illegible] probability. The top of the figure is north.

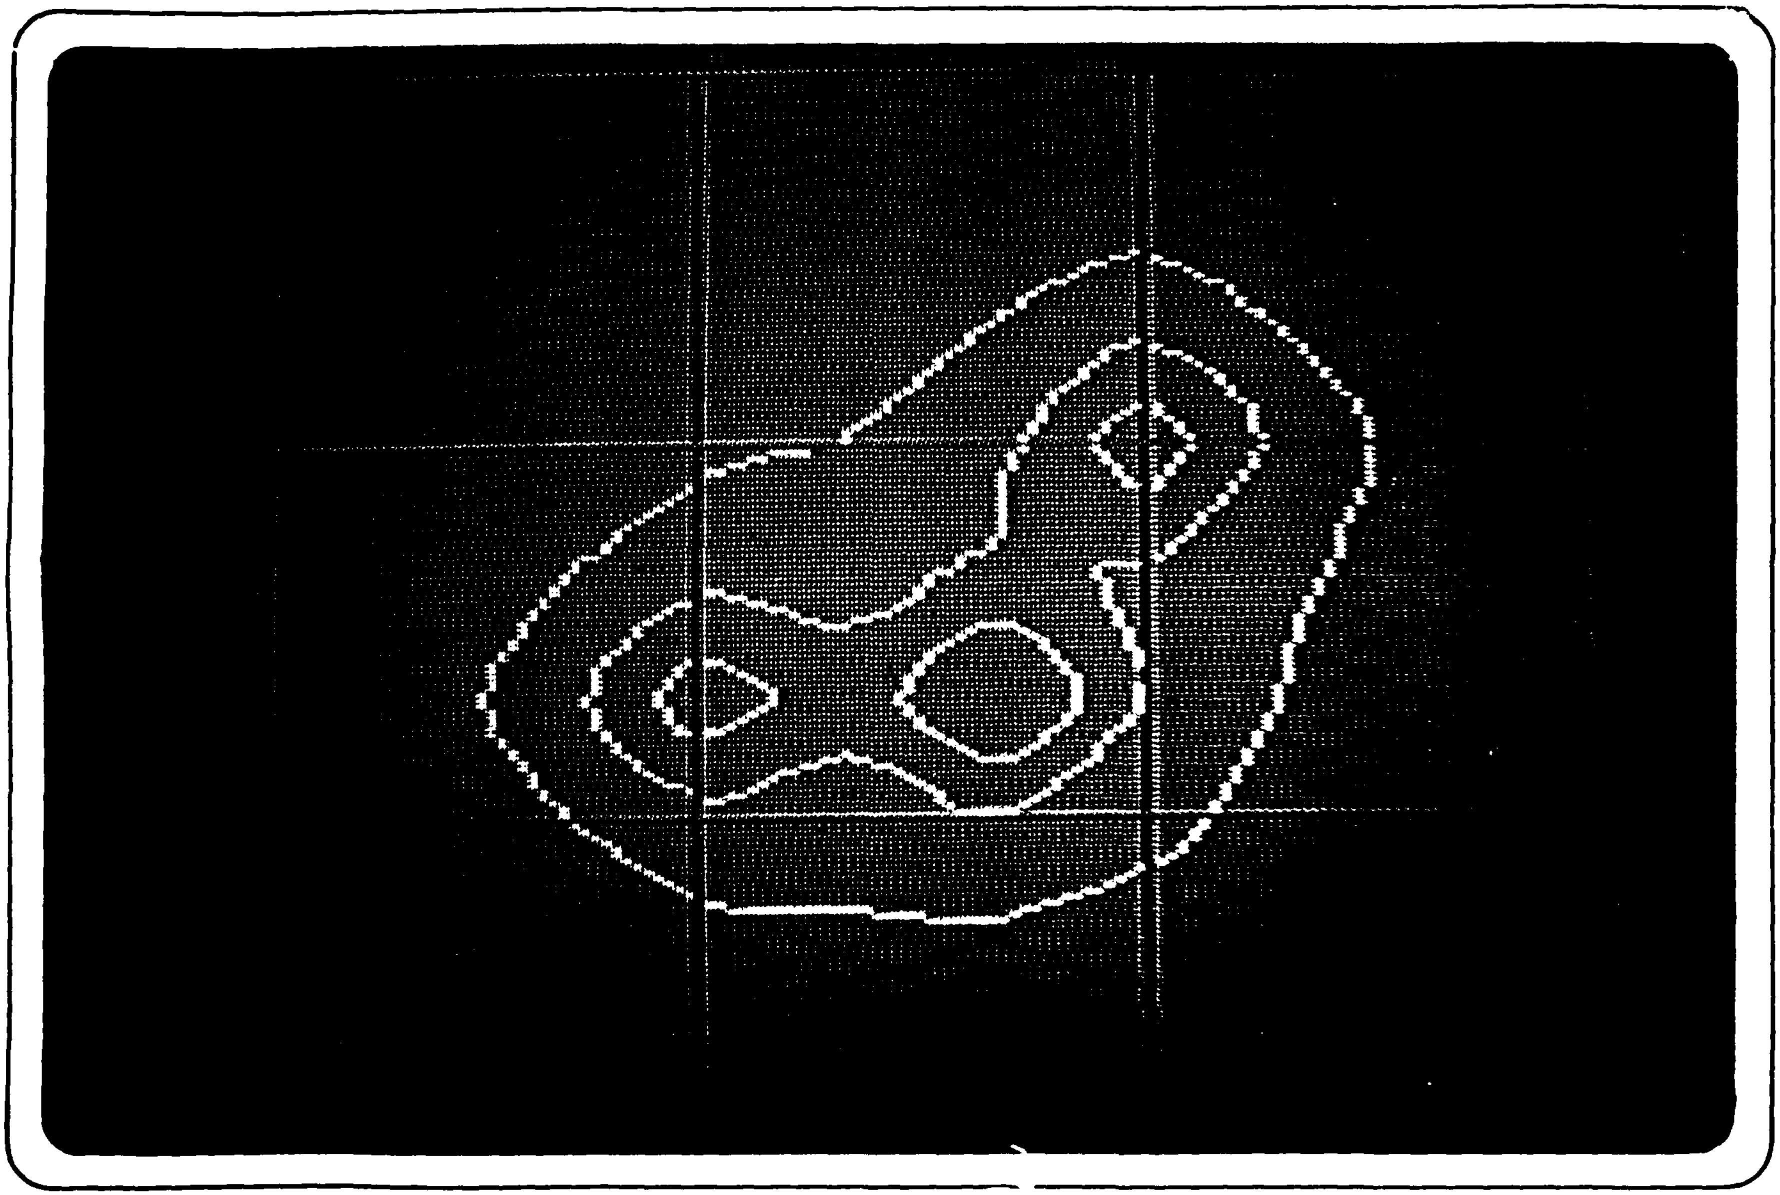

FIGURE 5
AN ALTERNATE DISPLAY OF THE PROBABILITIES OF AVOIDING DETECTION
FOR A GROUP OF THREE SHIPS

along the path and the speed with which the threat is moving (which determines how long the threat is located at each point along the path).

At a more intuitive level, regions with high detection rates correspond to areas with better radar coverage. If the detection rate in a region is doubled, then the probability that a threat can traverse the region undetected is halved. Thus, a display like the one shown in Figure 4 can be used to ensure that a balanced surveillance capability is established to protect important ships from an undetected threat. In particular, the display shows areas where radar coverage is relatively weak for certain approaches to high-value units. This type of display should be most useful for task force AAW decisions concerned with positioning defensive resources and controlling electronic emissions.

The primary assumption contained in this aid is that radar detection probabilities can be adequately modeled and calculated from detection rates. The detection rates are a function of the type of threat, the physical environment, electronic countermeasures and jamming used by both sides, and the status of various surveillance systems. These quantities and the corresponding detection rates can change with time. However, once the detection rate is determined, it is assumed to contain all the information necessary to compute the detection probabilities for a threat taking any path through the area. One consequence of this assumption is that the probability that a previously undetected threat will remain so depends only on the detection rate at its current location, and not on the route taken to reach that point. Another consequence is that the probability of detection is independent of whether the surveillance systems were able to detect a track at an earlier time. (This assumption can be removed by describing the performance of surveillance systems in probabilistic terms, and updating this

uncertainty as tracks are detected near the task force. However, this would require a difficult set of assessments about the uncertainty in sensor performance.).

The logic required to calculate the detection rates from physical data has not been developed as part of this project, and a separate research effort will be needed to specify appropriate algorithms. For the experimental aids described here and the display shown in Figure 4, the assumption was made that detection rates decline with the distance from each radar and that the detection rates for multiple radars are additive. This assumption was made to produce an illustrative set of displays, but is not required by any of the aids. The data used in all of the displays is hypothetical, and not intended to represent the capabilities of actual sensors.

The algorithms used to represent the capabilities of actual sensors should be based on empirical data describing the performance of each system under various operating conditions, including possible dependencies among radars. It may be possible to use theoretical models of radar performance to guide the specification of the algorithms, but the result should be a reasonably simple set of equations that correspond to the Navy's experience in using the sensors. Part of this development effort should be the specification of the external factors (e.g., weather, types of jamming, ECM) that have the most significant impact on detection probabilities. Once the algorithms have been specified, the data required to produce the displays will be available from existing AAW command and control systems. If the algorithms are developed to summarize the effectiveness of sensors, rather than model the details of electromagnetic propagation, the amount of computation needed to produce displays similar to Figure 4 should be well within the capabilities of fairly small computer systems.

Optimum and Likely Attack Routes

This aid produces two related displays: one showing the optimum routes for attacking aircraft or missiles trying to reach any point near the task force undetected, and a second showing regions containing likely, but not necessarily optimal, attack routes. These displays are generated from the detection rates described above, and from intelligence estimates of the enemy's objectives and knowledge of the task force's detection capabilities. The displays also reflect intelligence estimates of the direction from which the attack will be launched.

A typical display of the optimum routes to avoid detection is shown in Figure 6. This display shows the best attack routes when there are three radars generating the detection probabilities shown in Figure 4, the enemy is free to attack from any direction, his primary concern is avoiding detection, and he knows the location and performance of the radars. As Figure 6 demonstrates, the routes generated by this display need not be straight lines. The aid uses a dynamic programming algorithm to calculate the best route to each point in vicinity of the task force, and the algorithm is capable of diverting an attacker around a strong radar field in order to avoid detection. The algorithm is very efficient and converges quickly to a set of paths like those in Figure 6. It is described in detail in Section 2 of the Appendix.

Since the algorithm relies on the detection probabilities produced by the first decision aid, it requires the same assumption (i.e., that the performance of detection devices can be modeled by detection rates). In addition, the algorithm assumes that the detection rates are sufficiently smooth in the vicinity of the task force that flight paths can be approximated by paths consisting of straight line segments through a grid of fixed reference points. In practice, this is not a restrictive assumption and the computational

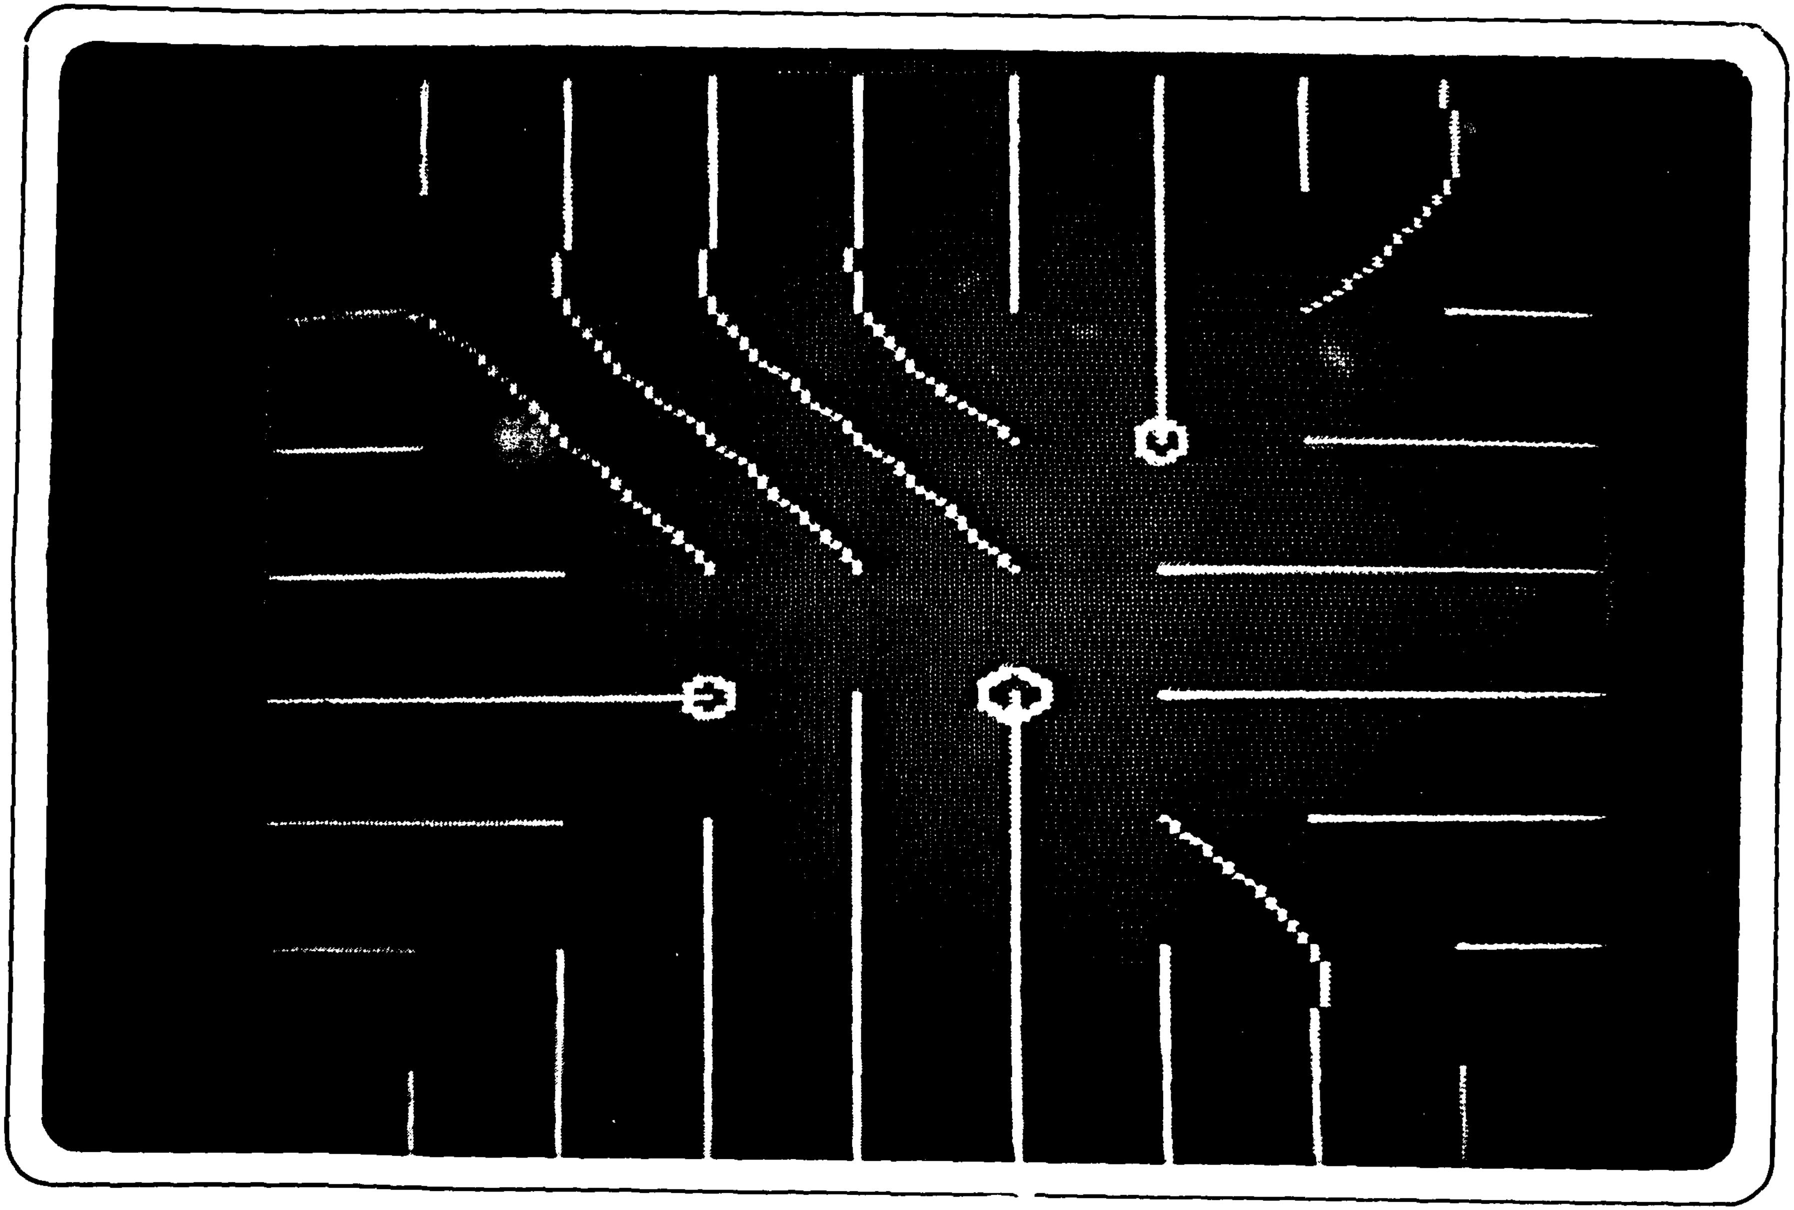

FIGURE 6
OPTIMUM ROUTES TO MINIMIZE DETECTION FOR AN ATTACK FROM ALL DIRECTIONS

errors caused by using it can be minimized by using a fine grid of reference points. The grid used to generate Figure 6 is fairly coarse (approximately 10 by 10), but the paths are very similar to those that would be generated using a finer grid.

The attack routes shown in Figure 6 change significantly if the intelligence estimates used to generate them are modified. For instance, Figure 7 shows the optimum routes when the attack is assumed to come from the north only (i.e., the top of Figure 7). In this case, the best way to reach many locations near the task force (including the central ship) is to take a circular route around the task force and attack from the east or west.

The display in Figure 7 uses the same assumptions as those that generated Figure 6, except for the direction of the attack. In particular, both figures assume that the enemy's primary objective is to minimize detection. If we change this assumption so that the enemy's objective is to minimize both the probability of detection and the distance traveled to reach any given point, then the optimal paths shown in Figure 7 (i.e., those for a northern attack) change to those in Figure 8. In Figure 8, the attack routes are much more direct, reflecting the need to minimize distance (i.e., conserve fuel) as well as avoid detection. The parameter that defines the relative importance to the enemy of minimizing detection and distance can be controlled by AAW personnel, based on intelligence estimates or observed enemy behavior. However it is unlikely that this parameter will have to be changed often.

The user also can control part of the aid's logic describing the enemy's knowledge of the task force's surveillance capability. If there is reason to believe that the enemy does not know the exact location of the radars, or is unaware that some of the detection

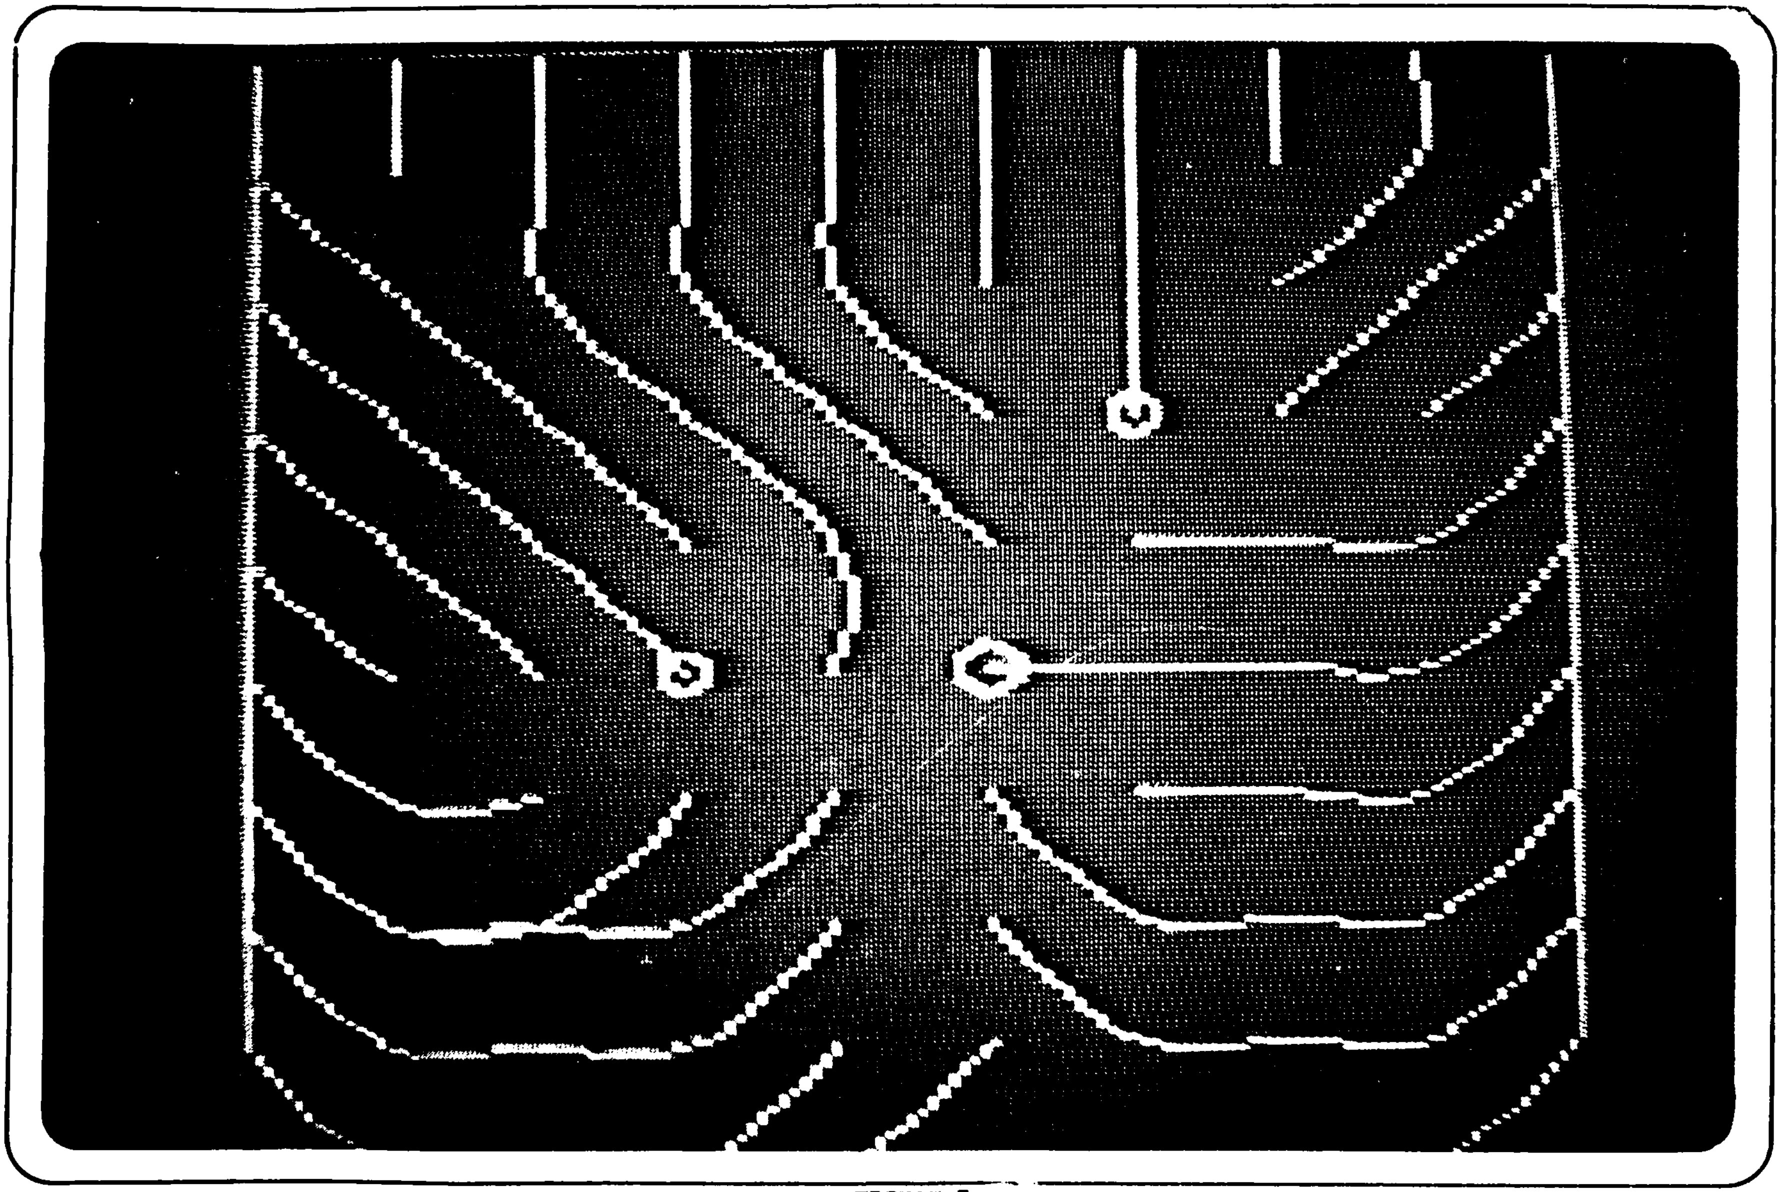

FIGURE 7

OPTIMUM ROUTES TO MINIMIZE DETECTION FOR AN ATTACK FROM THE NORTH

The top of the figure is north.

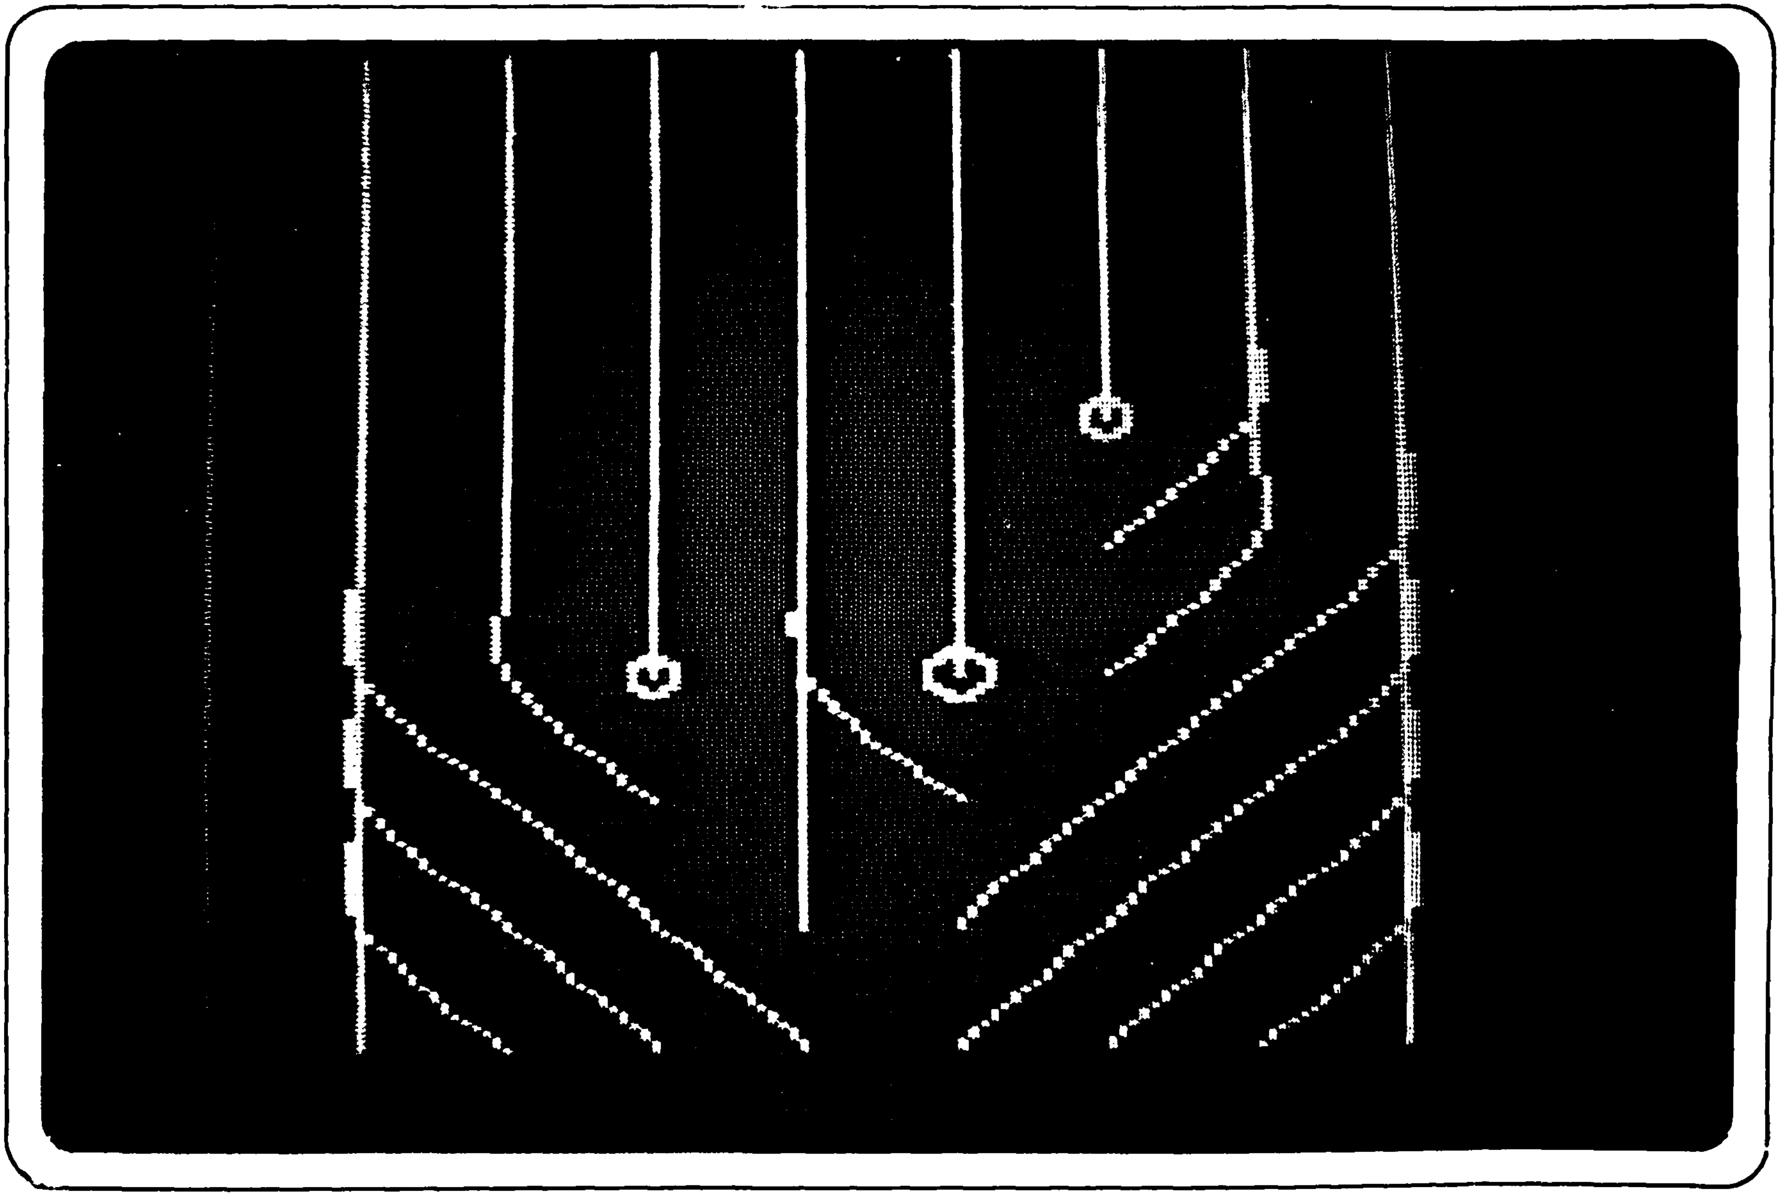

FIGURE 8
OPTIMUM ROUTES TO MINIMIZE DETECTION AND DISTANCE TRAVELED FOR AN ATTACK FROM THE NORTH
The top of the figure is north.

devices are operating (as in the case of passive detection devices), this information can be entered and the display of optimal attack routes will change accordingly.

It would be unwise to concentrate too many defenses on the optimal attack routes If the enemy's objectives and knowledge differ from our estimates, or the enemy decides to take a suboptimal route, defensive capability may be needed elsewhere. Therefore, the decision aid produces a second type of display showing the relative suboptimality of various paths to a given target. An algorithm that produces this display is described in detail in Section 3 of the Appendix. If we assume that an enemy is more likely to take routes that minimize the probability of detection or the distance traveled, then this display indicates the relative likelihood that the enemy will select routes that pass through various regions. This interpretation is consistent with the mixed (i.e., probabilistic) strategies that would result from a game theoretic approach to the problem.

Figure 9 shows this display for the case where the enemy is free to attack the central ship from any direction, his primary concern is avoiding detection, and he knows the location and performance of our radars. In this case the most likely attack routes (i.e., those that do the best job of minimizing detection) are from the south and east. Figure 10 shows the same display when the attack is assumed to come from the north. This display shows that an enemy attempting to minimize detection is likely to take an indirect route and approach the target (the central ship) from the east.

The displays produced by this aid do not require physical data other than that needed to generate the detection rate probabilities. However, the calculation of optimal and likely attack routes requires several important assumptions about the enemy's tactics, objectives,

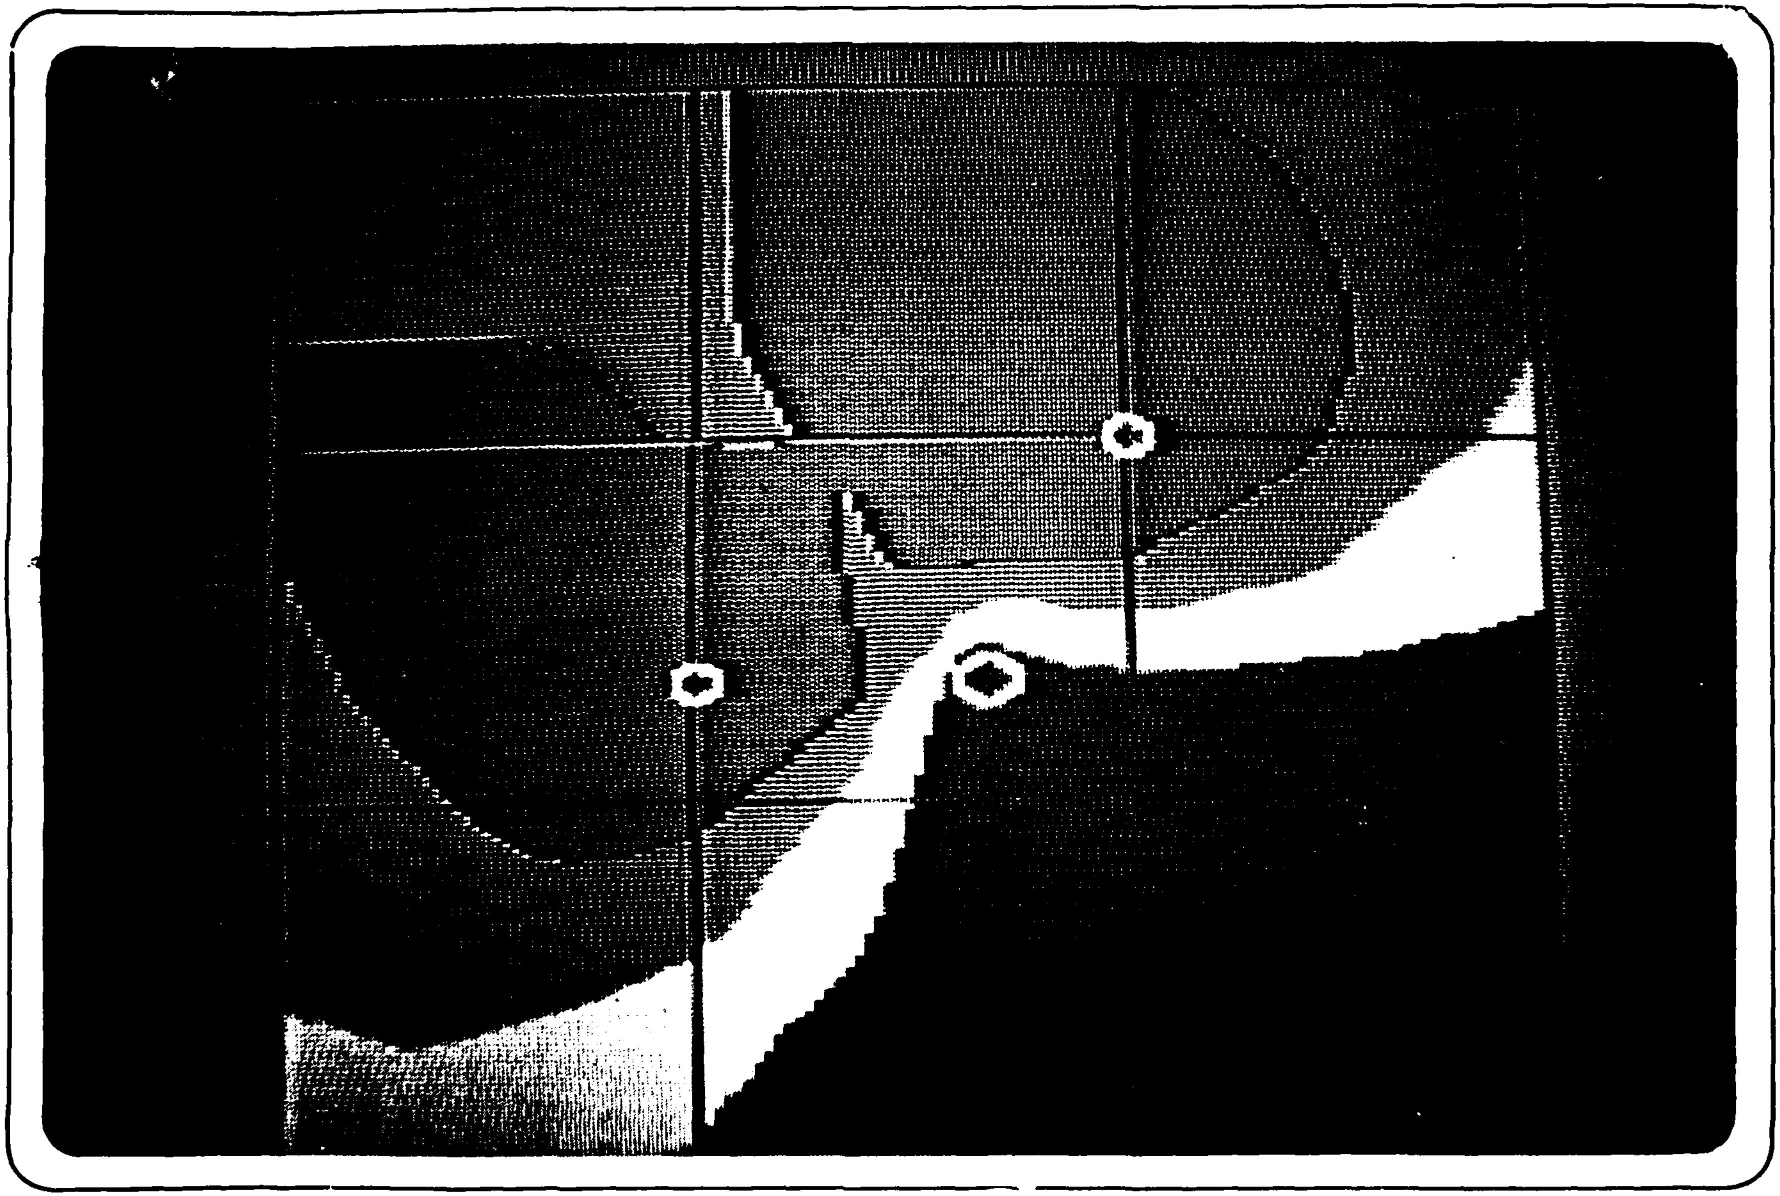

FIGURE 9

THE LIKELIHOOD THAT AN ATTACK WILL PASS THROUGH ANY POINT NEAR THE TASK FORCE TO REACH THE CENTRAL SHIP FOR AN ATTACK FROM ALL DIRECTIONS

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

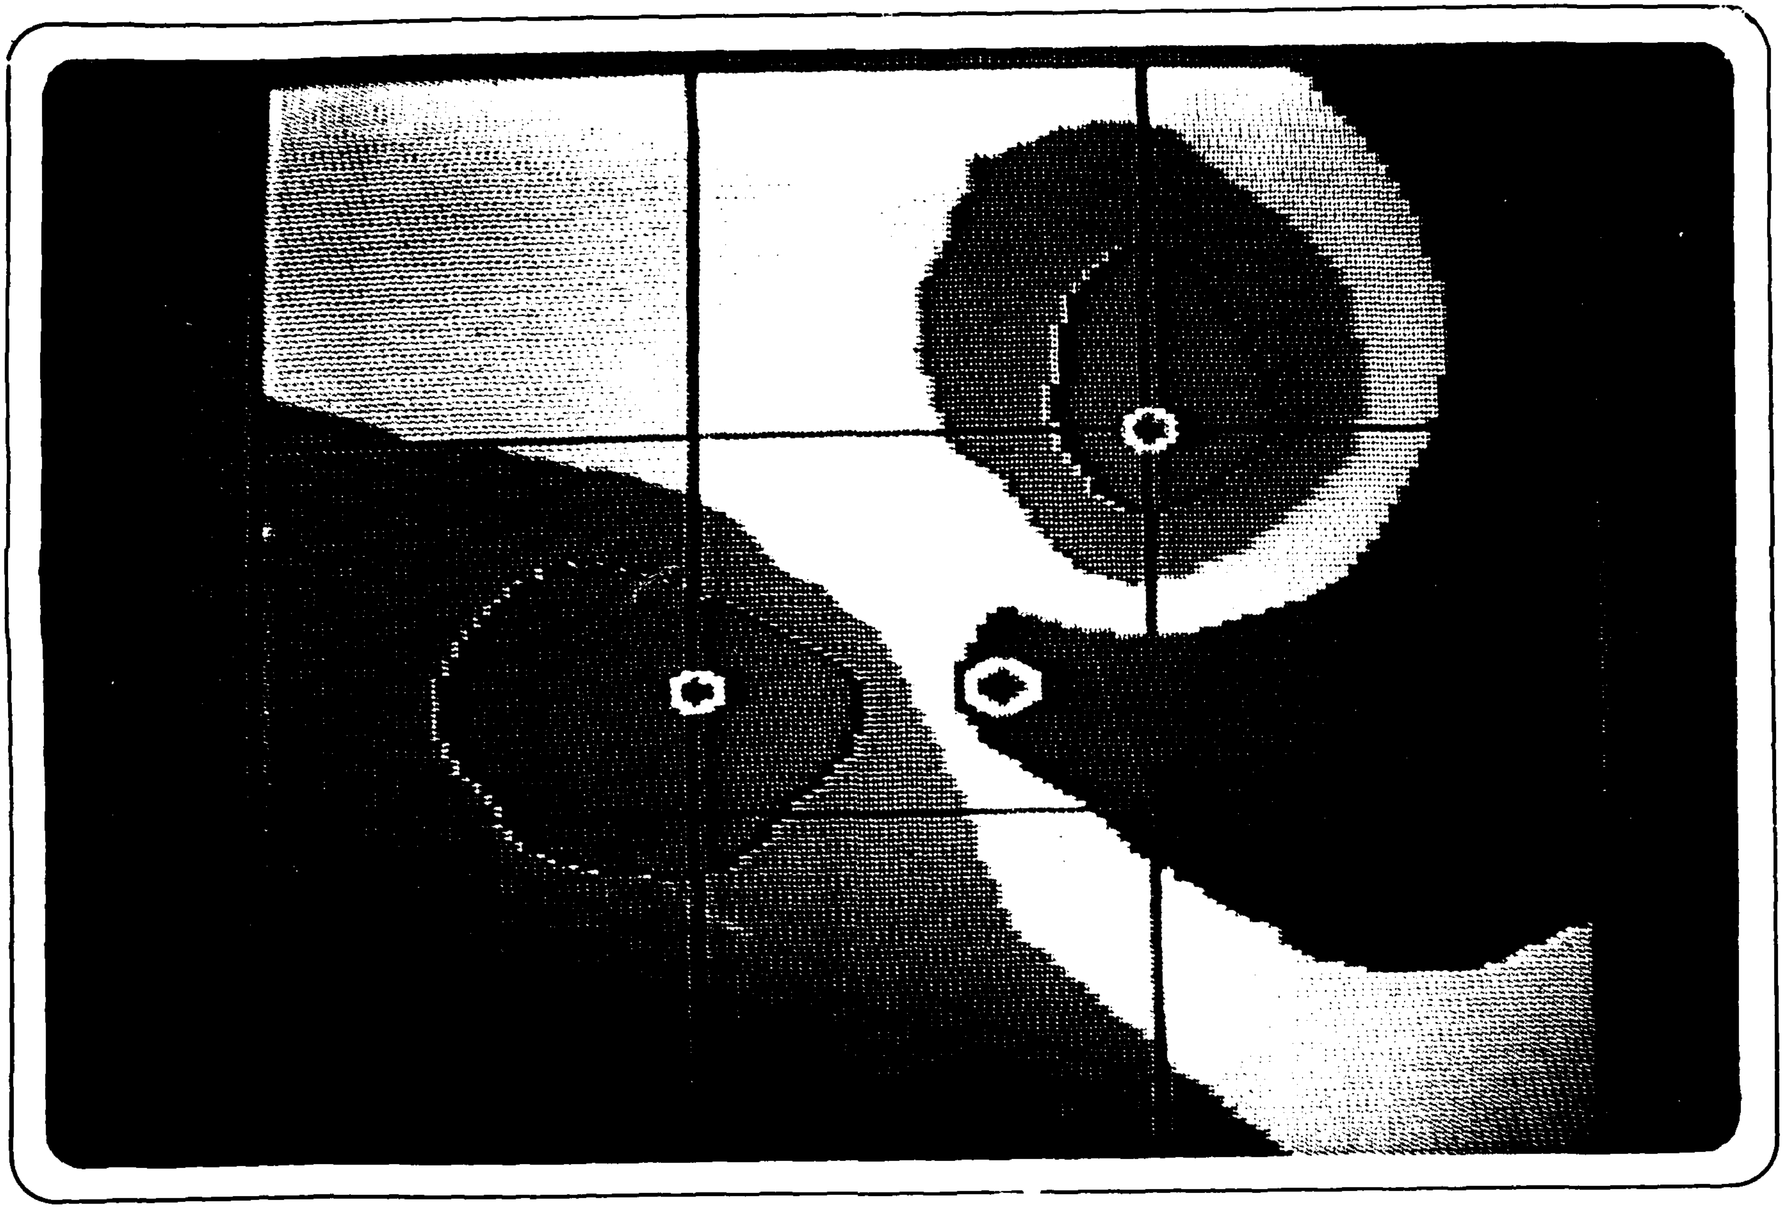

FIGURE 10
THE LIKELIHOOD THAT AN ATTACK WILL PASS THROUGH ANY POINT NEAR THE TASK FORCE TO REACH THE CENTRAL SHIP FOR AN ATTACK FROM THE NORTH
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

and knowledge. These estimates should be supplied by intelligence analysts, with the aid serving to show the implications of their assessments for allocating defensive resources.

Surveillance Penetration Probabilities

This decision aid uses information produced by both aids described above to produce a display of the maximum probability that a threat can penetrate undetected to any point in the vicinity of the task force. This display is based on a "worst case" analysis using the optimum attack routes. Figure 11 shows the surveillance penetration probabilities for the case where the detection field is that shown in Figure 4, the enemy can attack from any direction, he is trying to avoid detection, and he knows the location and performance of the radars. In this case, the enemy follows the optimal paths shown in Figure 6, and the probability that he will reach any point near the task force is represented by the colored regions in Figure 11. The boundaries of these regions correspond to various probabilities of penetration, and those boundary values can be selected by the user. The probabilities of penetration decline as the threat gets closer to the radars. (The exact values of the probabilities corresponding to the contours in Figure 11 are not significant since the data used to generate this display is hypothetical.)

Like the detection probabilities in Figure 4, the display in Figure 11 shows the implications of positioning surveillance systems and establishing an emissions control policy. Using this aid, AAW personnel could examine the detection capabilities associated with alternative allocations of surveillance systems and emissions control strategies, and select the one which provides the best mixture of good coverage and low visibility. For example, if a fourth radar is positioned southeast of the central ship, the surveillance penetration probabilities change to those shown in Figure 12. In

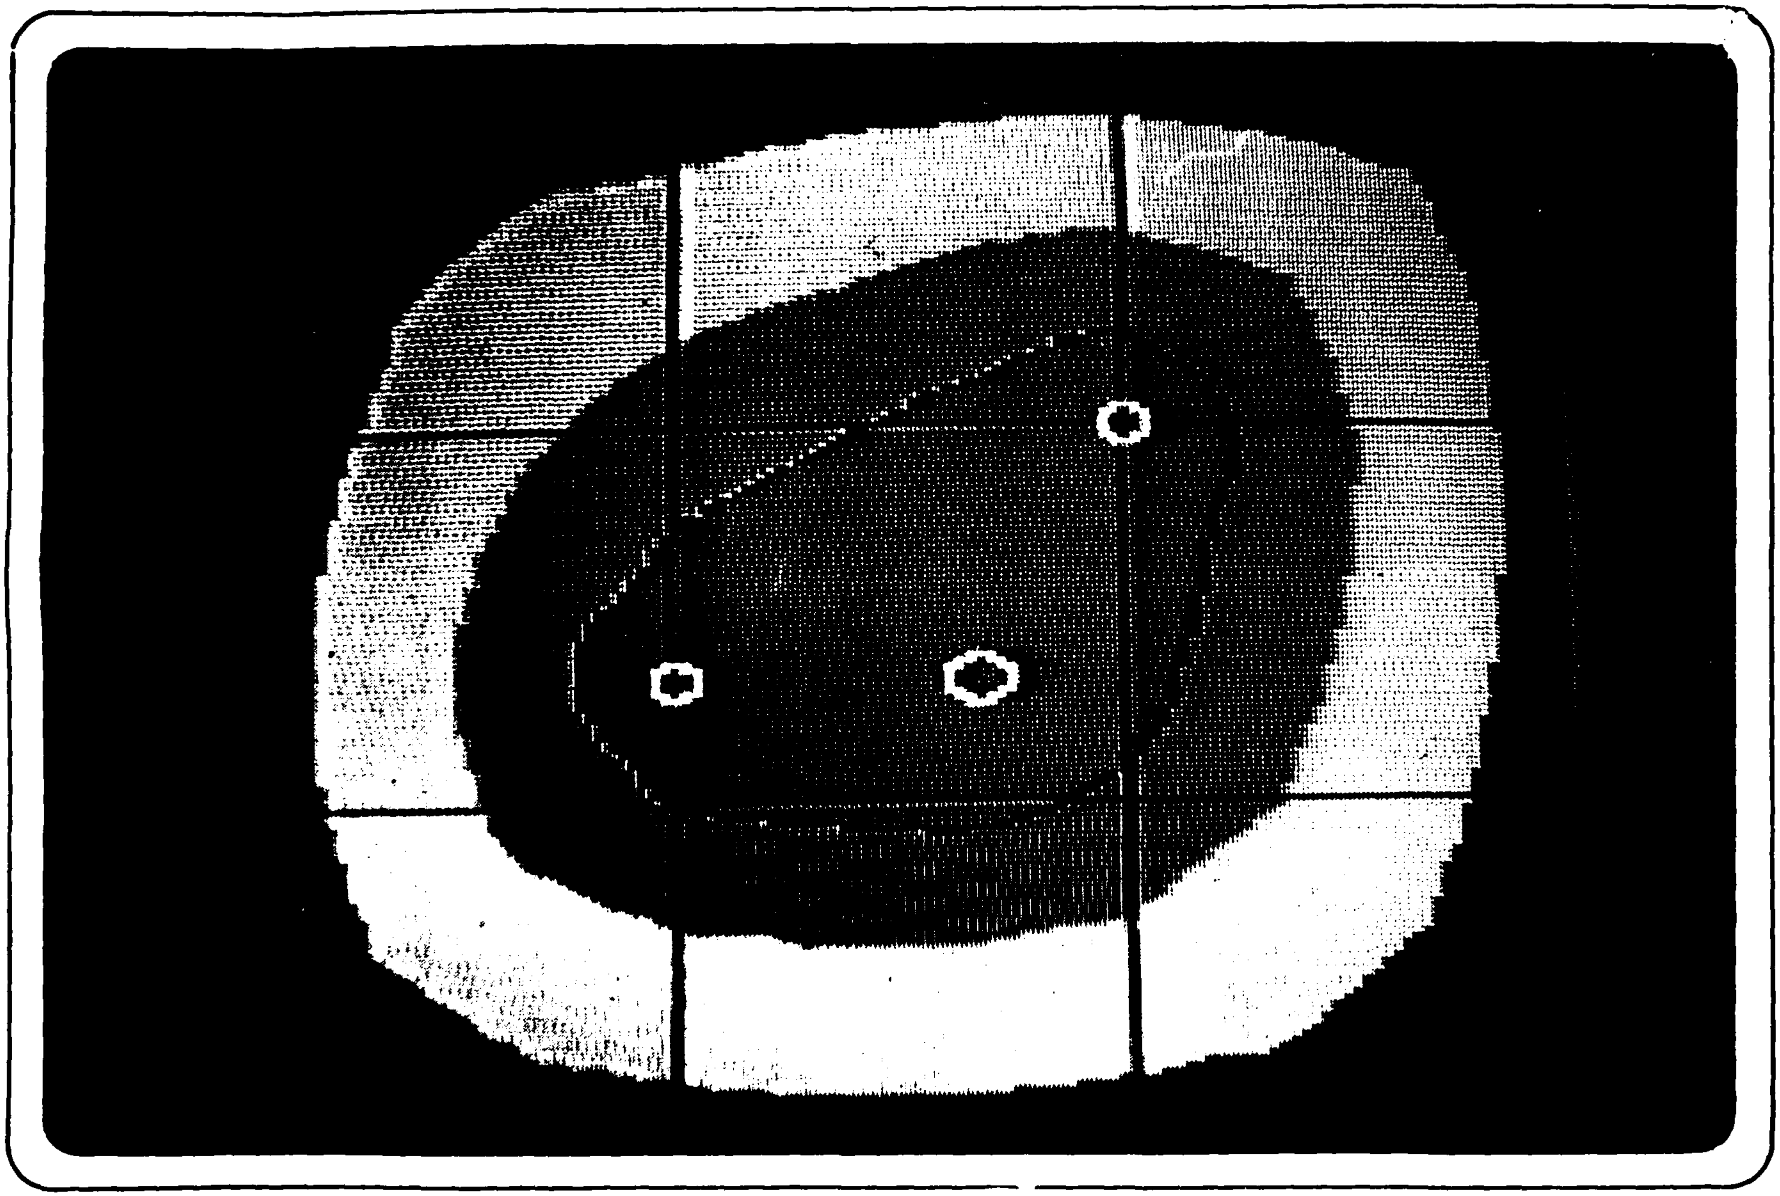

FIGURE 11
SURVEILLANCE PENETRATION PROBABILITIES FOR AN ATTACK FROM ALL DIRECTIONS
Red corresponds to regions of highest probability, and blue c[illegible]nds to lowest probability.

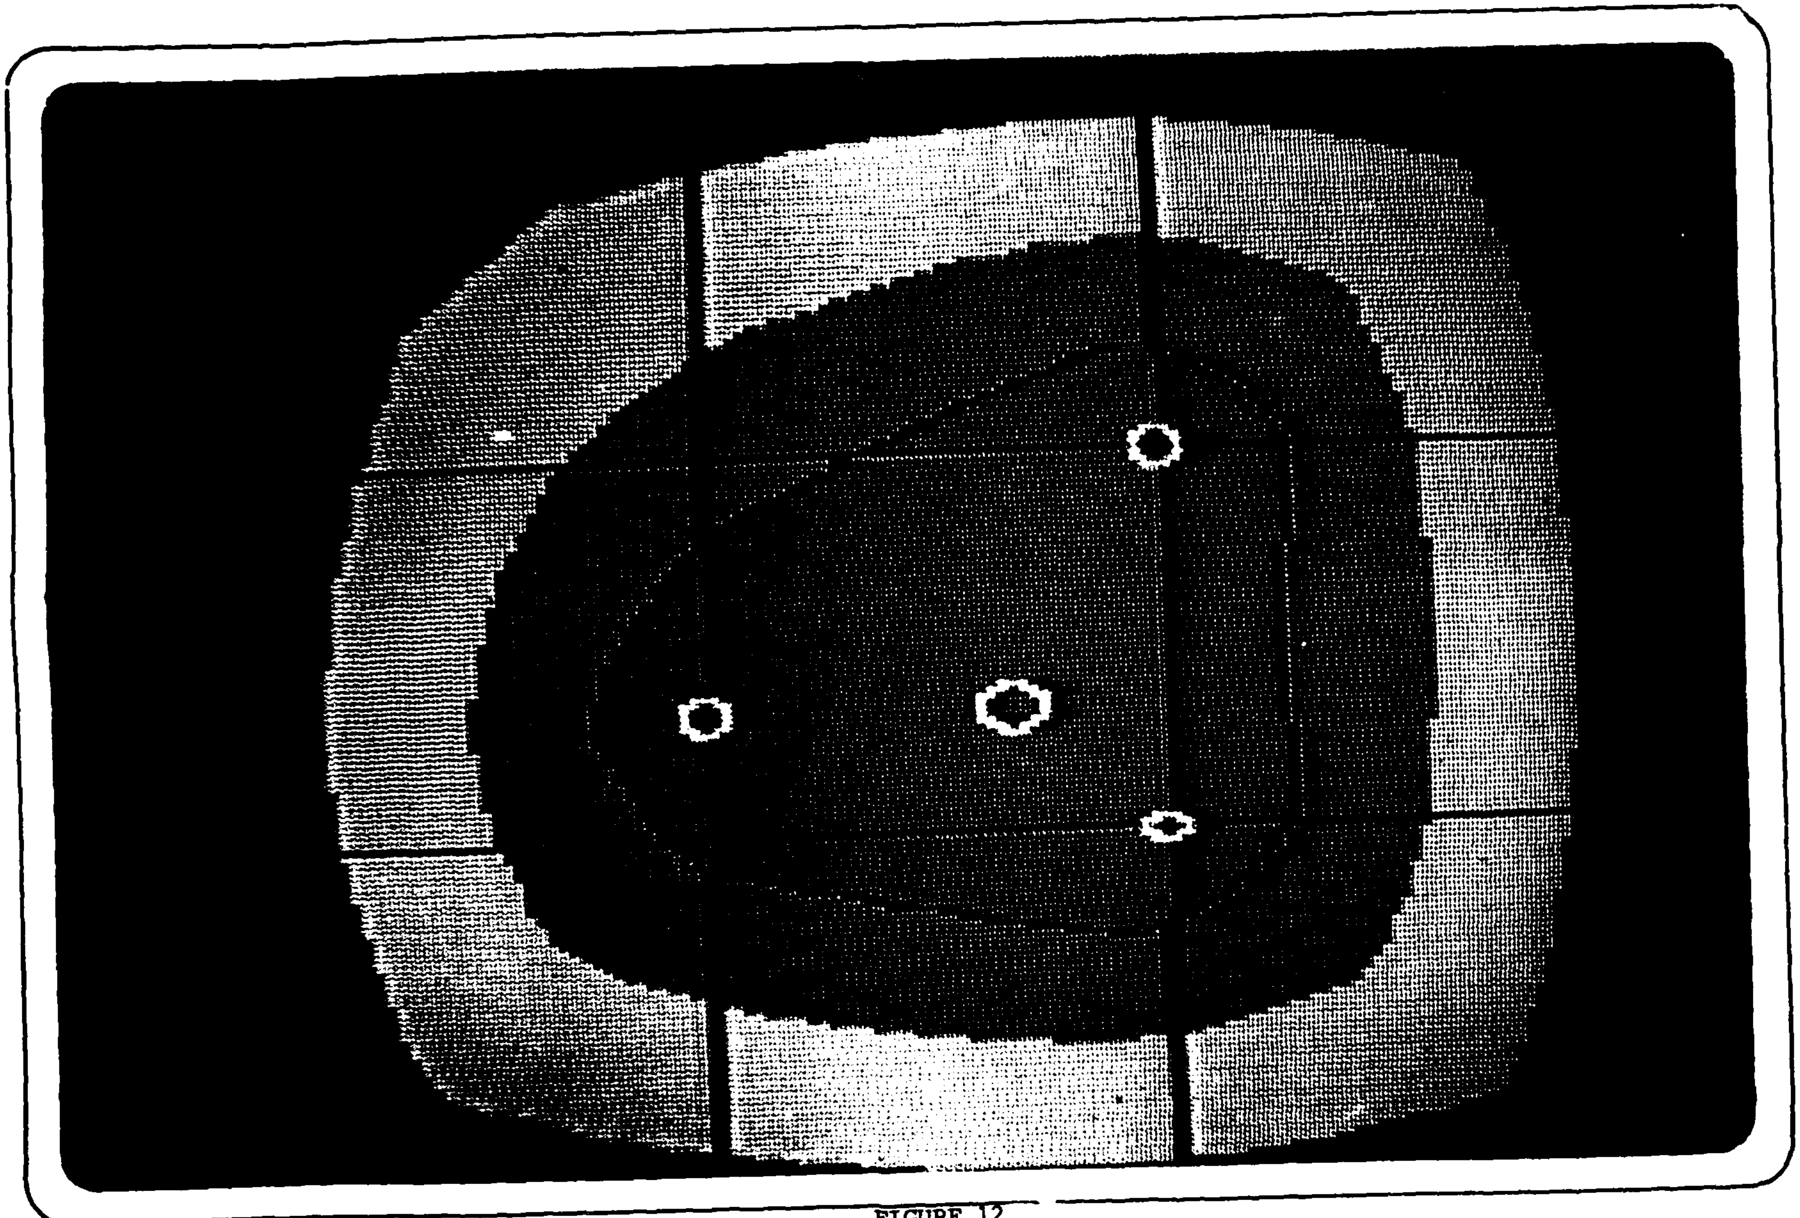

FIGURE 12

SURVEILLANCE PENETRATION PROBABILITIES FOR AN ATTACK FROM ALL DIRECTIONS
WHEN A FOURTH RADAR IS ADDED

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

this case, the inner region of low penetration probabilities has been expanded by the addition of the fourth radar. Even though Figure 12 contains more individual units than Figure 11, the amount of information conveyed to the user by the probability regions remains approximately the same. The addition of many more ships and radars would still produce a display with approximately the same amount of information for AAW personnel to interpret.

Although the detection probabilities in Figure 4 and the surveillance penetration probabilites in Figure 11 both show the effects of the same surveillance capability, there is a major difference between the two displays. Figure 11 also shows the implications of the assumptions about enemy tactics, objectives, and knowledge. For example, if it is assumed that the attack will come from the north, the display of surveillance penetration probabilities in Figure 11 changes to that shown in Figure 13. In this case, the region of low penetration probabilities has expanded to include most of the area south of the task force since a threat entering from the north would have to travel a long indirect route to reach this area. (Figure 7 shows the minimum detection routes for this case.)

If it assumed that threats attacking from the north will attempt to minimize both the probability of detection and the distance traveled to reach the target, the penetration probabilities in Figure 13 change to those in Figure 14. In Figure 14, the region of low penetration probabilities has expanded to the east and west relative to the display in Figure 13, to reflect the assessment that the enemy is unwilling to travel the long distance required to pass through these locations on the way to the task force.

In addition to the surveillance penetration probabilities for threats that have never been detected, the decision aid can display them for a lost track. Figure 15 shows the probability that a track lost north of the task force (i.e., at the blue dot) can

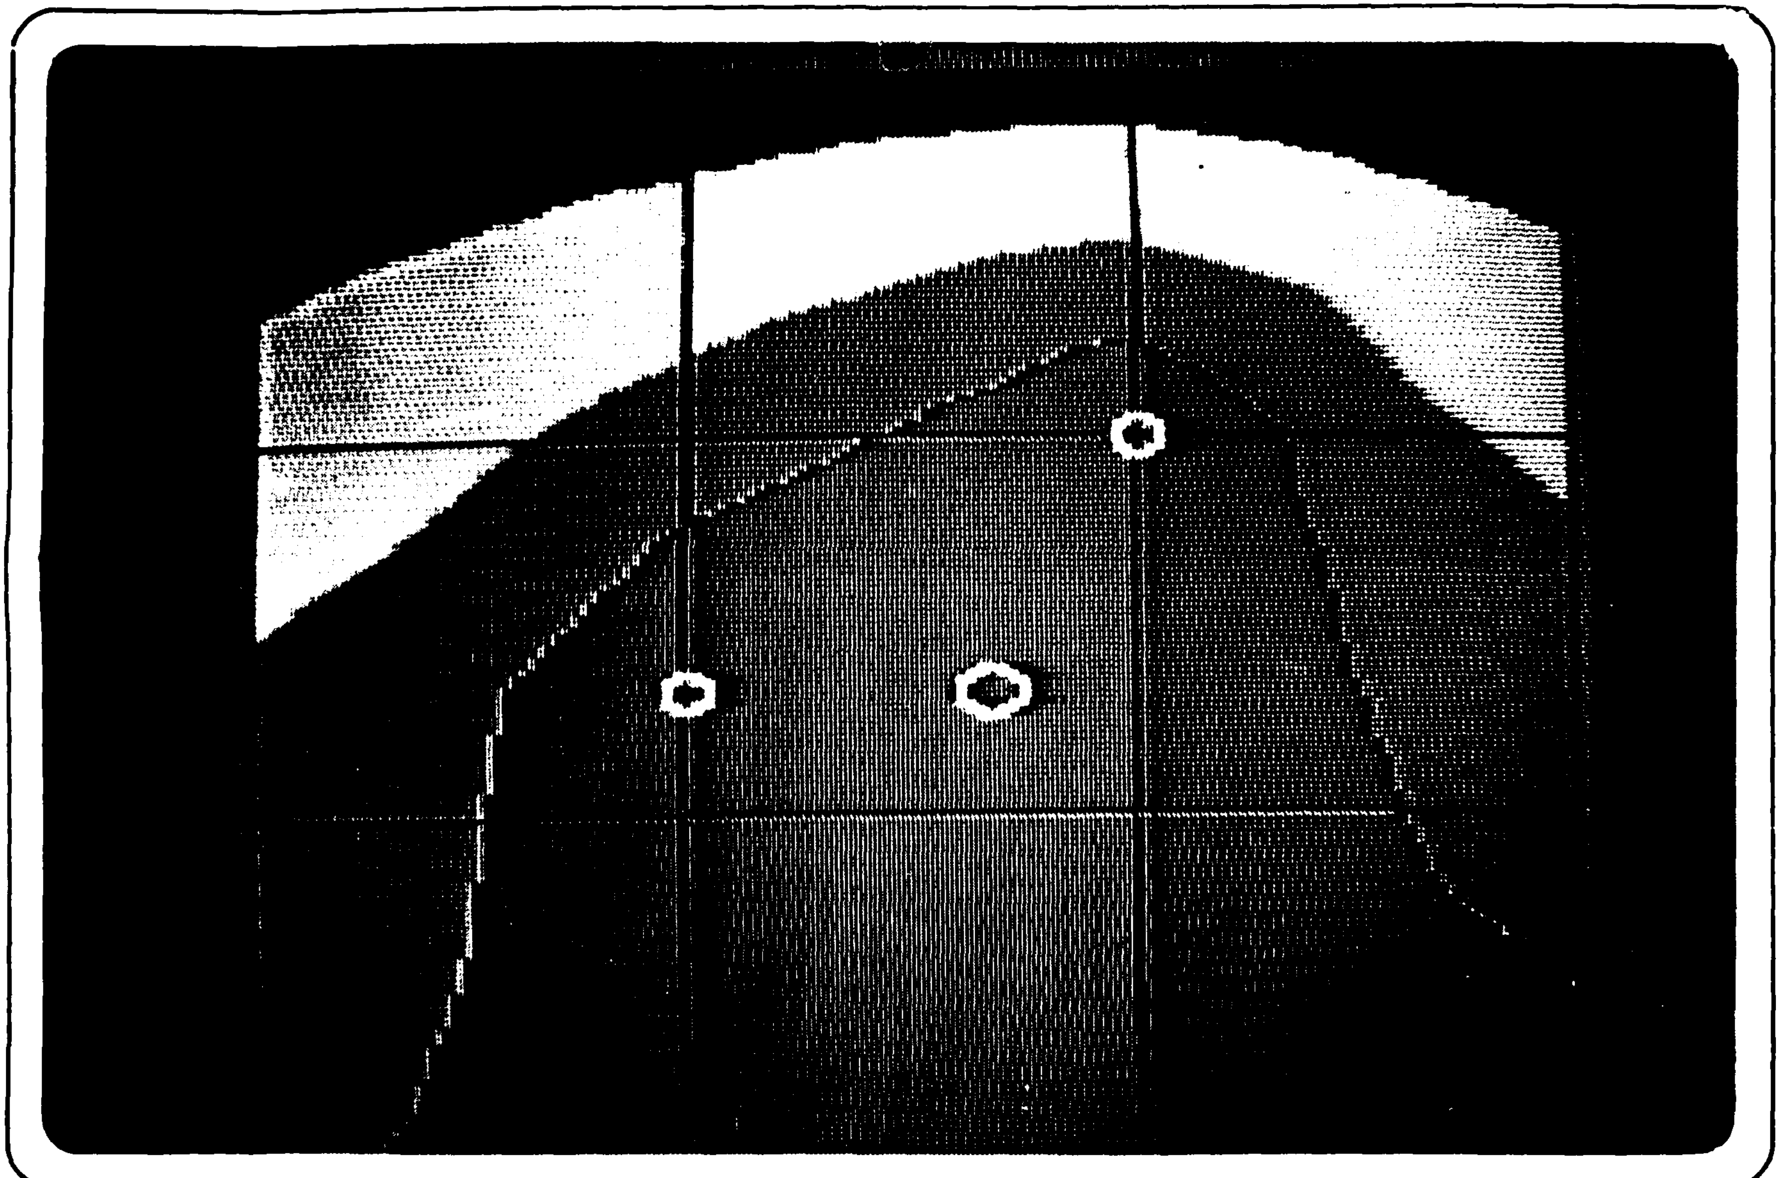

FIGURE 13

SURVEILLANCE PENETRATION PROBABILITIES FOR AN ATTACK FROM THE NORTH

Red corresponds to regions of highest probability, and blue corresponds to lowest probability. The top of the figure is north.

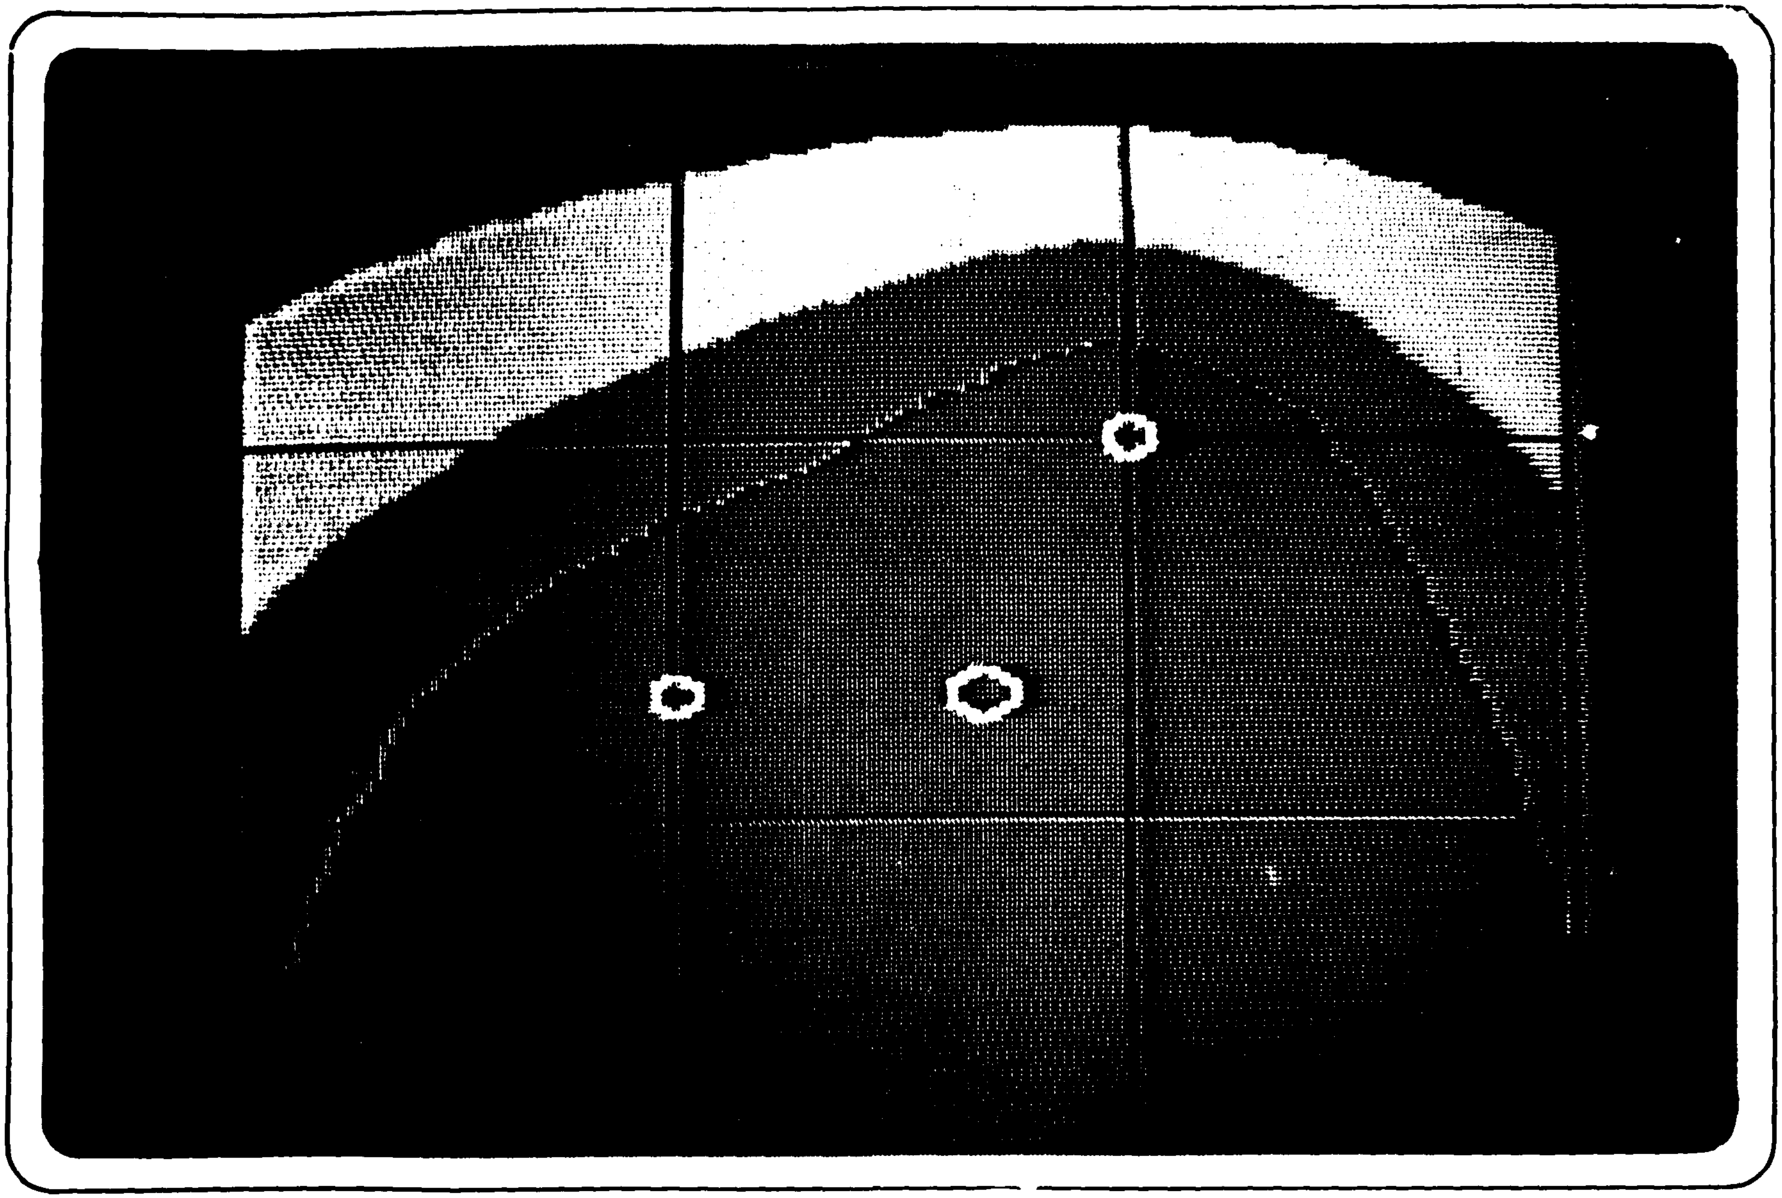

FIGURE 14
SURVEILLANCE PENETRATION PROBABILITIES FOR AN ATTACK FROM THE NORTH
WHEN THE ENEMY'S OBJECTIVE IS TO MINIMIZE DETECTION AND DISTANCE TRAVELED
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

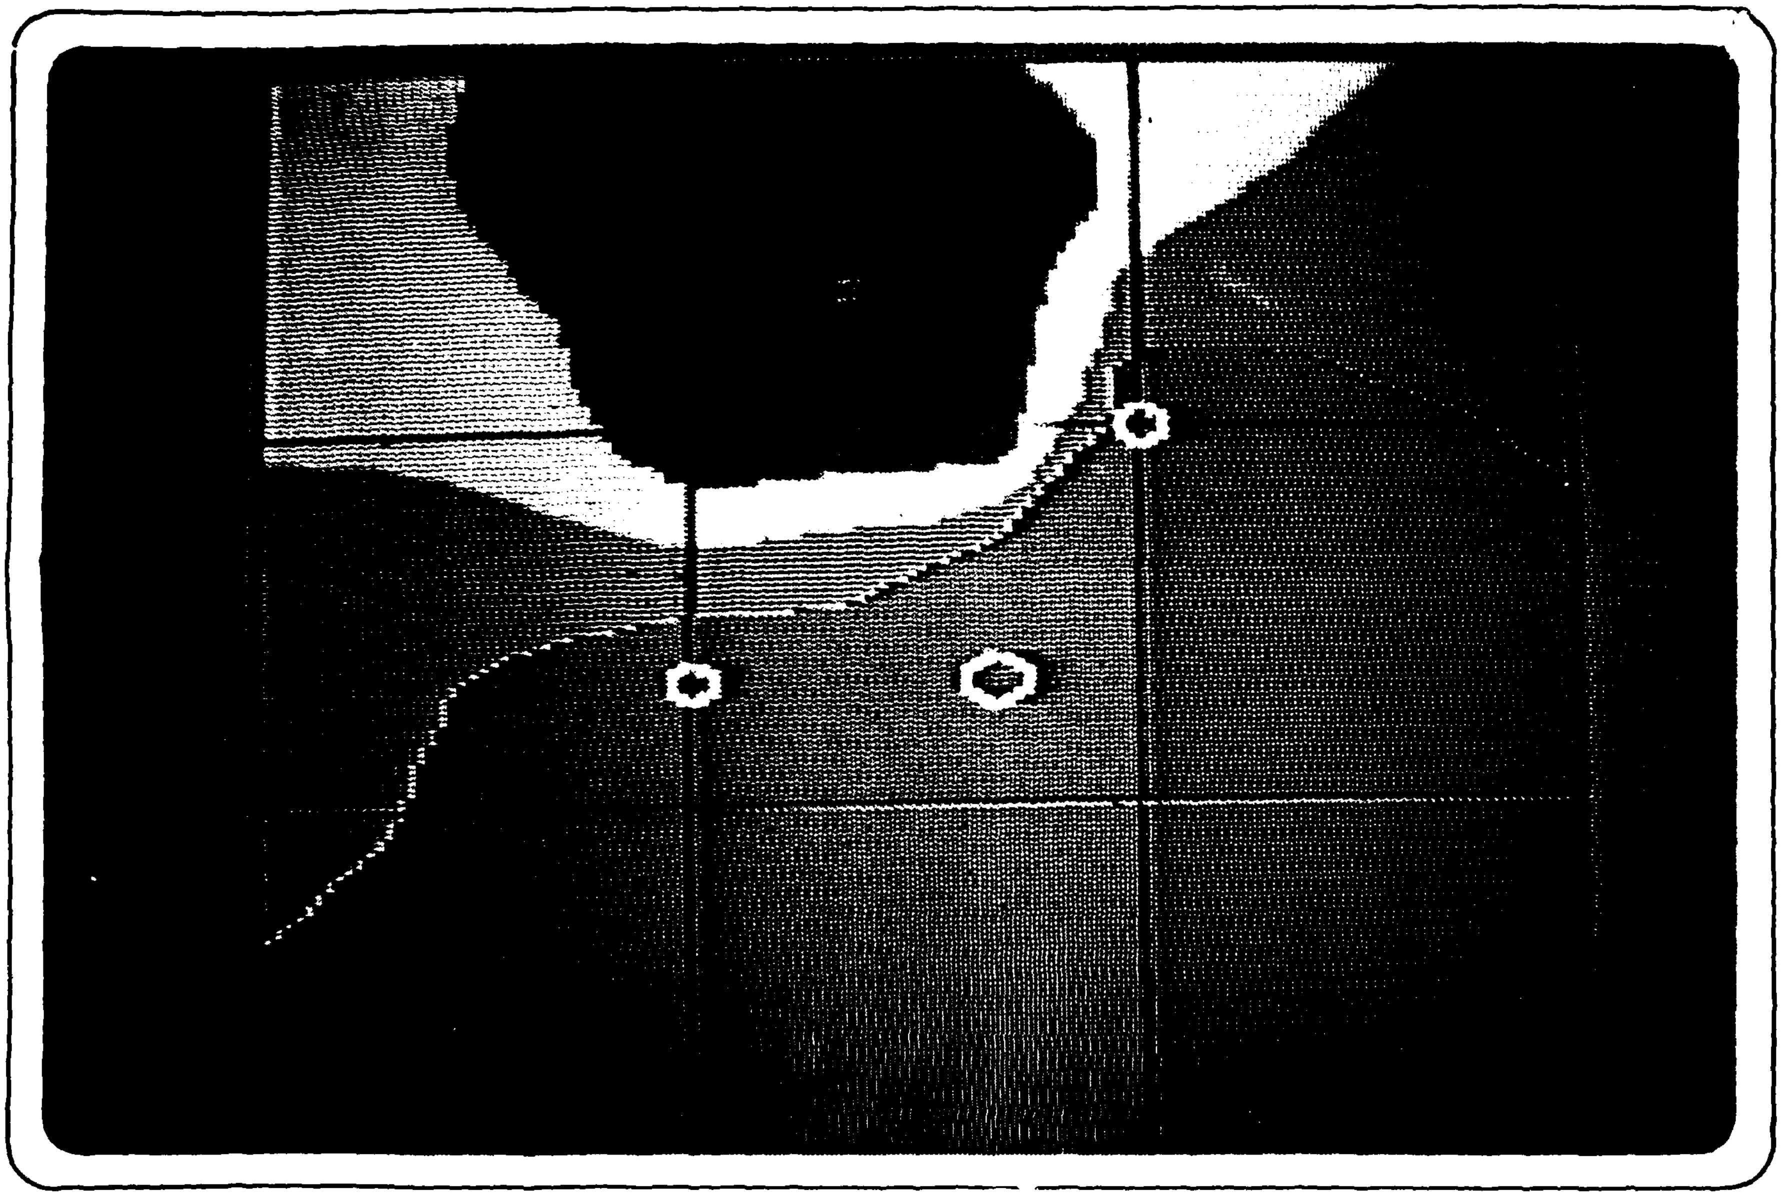

FIGURE 15
SURVEILLANCE PENETRATION PROBABILITIES FOR A LOST TRACK
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

penetrate undetected to any other location. This display shows the surveillance penetration probabilities immediately after the track is lost, based on the assumption that the threat is primarily interested in avoiding detection and knows the location and performance of the radars. Unlike the previous displays, this one is specific to a single track and should only be generated for AAW personnel processing that track. It is possible to combine this display with the ones showing surveillance penetration probabilities for tracks that have never been detected and show the probability that any track (including the lost track) can reach various locations undetected. However, it is more meaningful to combine the location and first-detection probabilities produced by the following decision aids.

The displays of surveillance penetration probabilities do not require any additional data or assumptions beyond those required to calculate detection probabilities and likely attack routes, except for information about lost tracks. This information is readily available from existing AAW systems. Furthermore, very little additional data processing is needed to calculate and display the surveillance penetration probabilities once the detection probabilities and attack routes have been determined.

Location Probabilities

This decision aid is closely related to the previous one, but it produces a display of current, rather than projected, threat status. The display shows the probability that one or more undetected threats are currently located at any point in the vicinity of the task force. The display can be generated for threats that have never been detected and for lost tracks, but different assumptions and data are needed for these situations. A complete technical description of this aid is presented in Section 4 of the Appendix.

Since this decision aid deals with the current location of threats, it is relevant to AAW personnel at the unit level who are responsible for detecting and tracking threats, and identifying tracks. In addition, it could provide support for decision-makers at the task force level who must position their forces to defend against current threats.

A typical display of location probabilities for a lost track is shown in Figure 16. In this display, a track was lost at a location north of the task force and it remained undetected for 10 minutes. The colored regions show the probability that the track is currently located at any point within 10 minutes flying time of the point at which it was lost. This display is based on the assumption that the track is equally likely to travel to any point it can reach in 10 minutes. In other words, the prior probability distribution for the location of the track is assumed to be uniform. (This assumption will be changed later.) Since the track has not been detected and the chances of detection are higher as it approaches the task force, the track is probably moving away from the task force into regions where detection is less likely. This effect is illustrated in Figure 16.

The display in Figure 16 is generated by calculating the probability that a lost track will be detected again if it travels along the minimum detection routes (or the routes corresponding to a combination of minimum detection and distance) from the point where it was lost to any point it can reach in the time available. Bayes' Rule is used to combine this set of probabilities with the prior probability distribution for the track's location to produce an updated distribution of location probabilities. The display shows the updated distribution.

If a track remains undetected for longer periods of time, the updated location probabilities change to reflect both the ability of the track to travel farther and the inability of the radars to regain

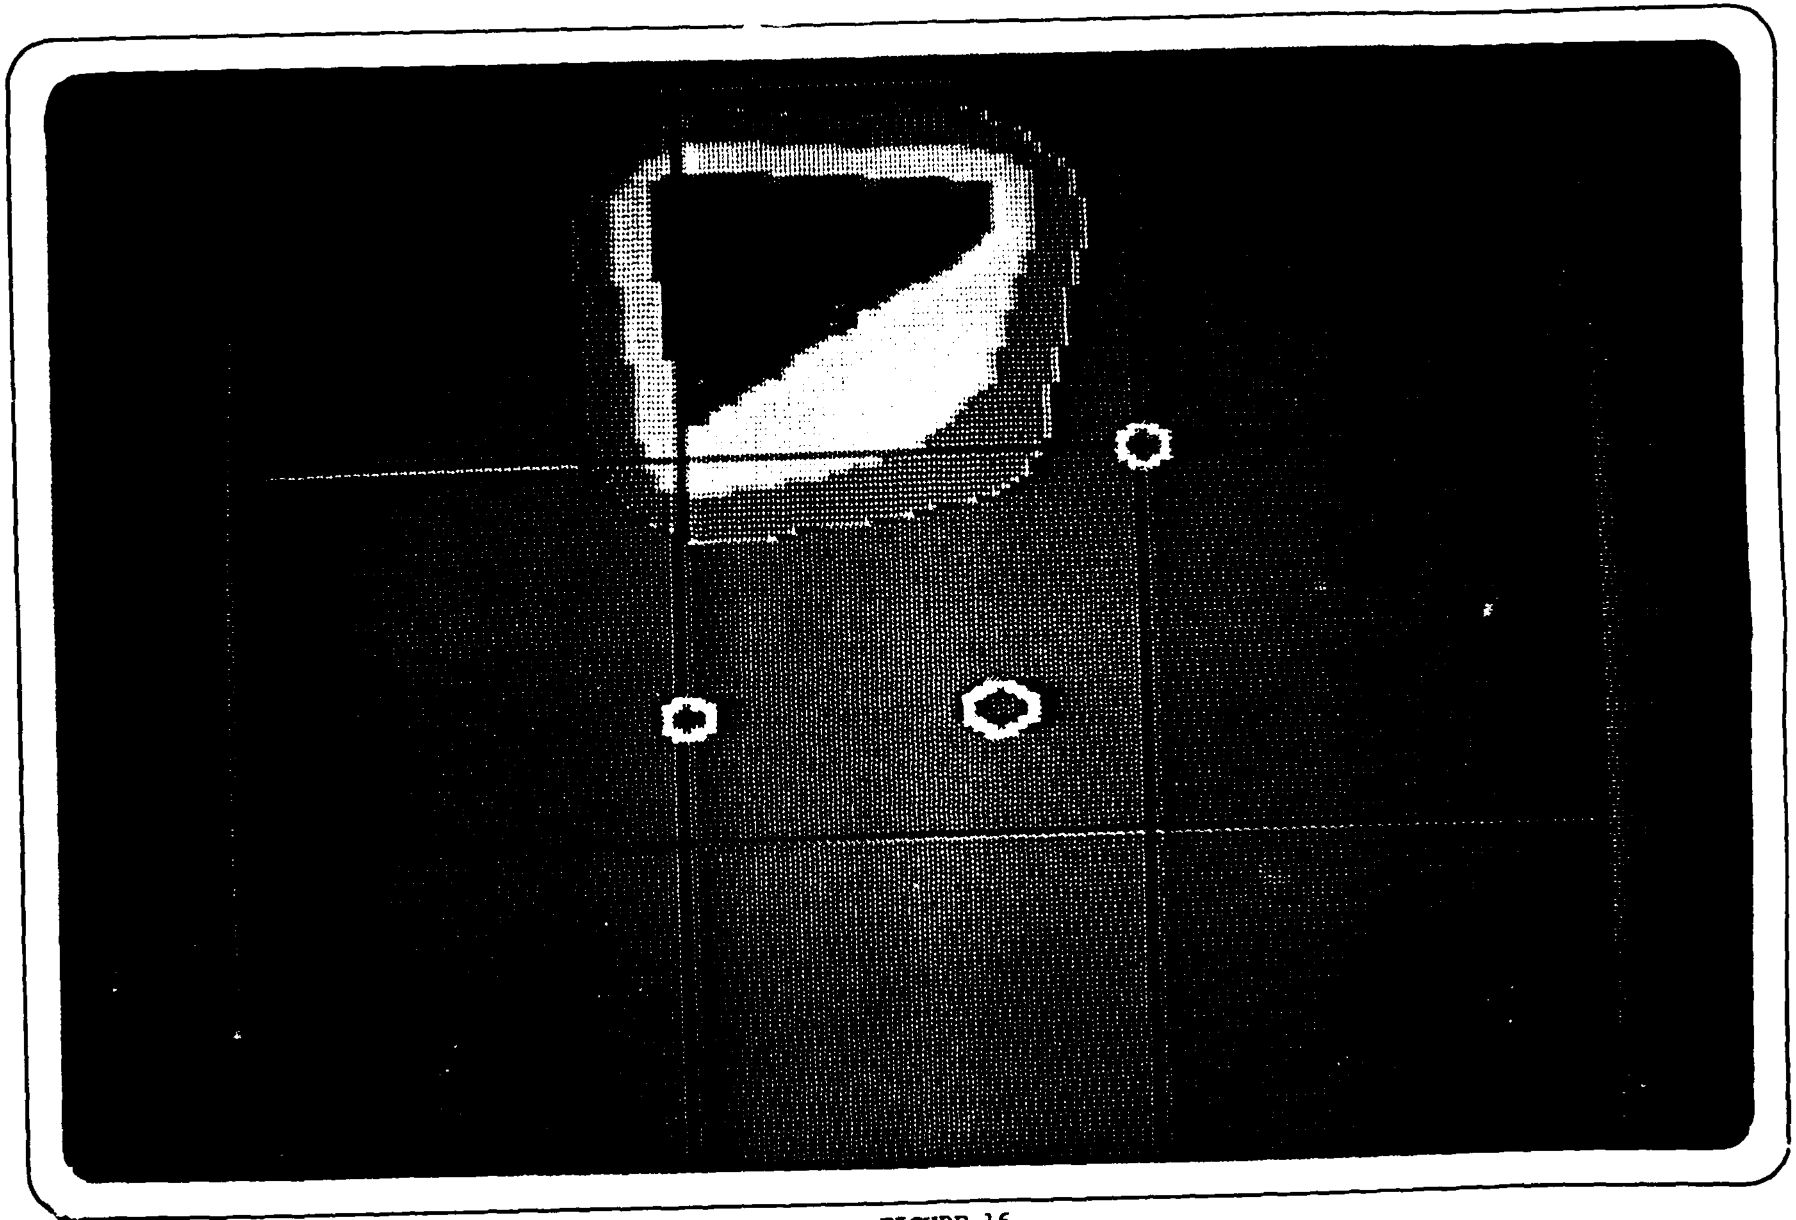

FIGURE 16
LOCATION PROBABILITIES FOR A LOST TRACK THAT REMAINS UNDETECTED FOR 10 MINUTES
GIVEN A UNIFORM PRIOR DISTRIBUTION FOR ITS LOCATION
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

contact. For example, if the lost track in Figure 16 remains undetected for an additional 10 minutes (i.e., it has been lost for 20 minutes), the updated location probabilities become those in Figure 17. In this case, the region likely to contain the lost track has expanded and moved farther from the task force.

The display of location probabilities uses data produced by all of the decision aids described earlier, and thus requires the same assumptions and input data. In addition, the algorithm for calculating updated location probabilities must have a prior probability distribution for the location of a lost track. This information could be assessed directly by AAW personnel, but it is doubtful that they will have time to do so. Alternatively, the prior distribution can be generated using a set of heuristic rules about the objectives and performance of air threats, and AAW personnel can accept this prior or modify it to reflect any special information or circumstances.

For instance, if we assume that a lost track is most likely to continue moving in the direction it was traveling at the time it was lost, the aid can automatically generate a prior distribution like that shown in Figure 18. In this example, the track was traveling south when it was lost at a point north of the task force (at the blue dot). The prior location probabilities are assumed to be highest for points south of track's initial position, and they decline with the extent of the course change required to reach them. Using this prior distribution produces the updated location probabilities shown in Figure 19, given that the track has avoided detection for 10 minutes. A comparison of Figures 16 and 19 shows that the prior distribution based on the track's direction at the time it was lost has significantly altered the updated location probabilities. Figure 19 indicates that the track is likely to be closer to task force than it was when it was lost (in keeping with

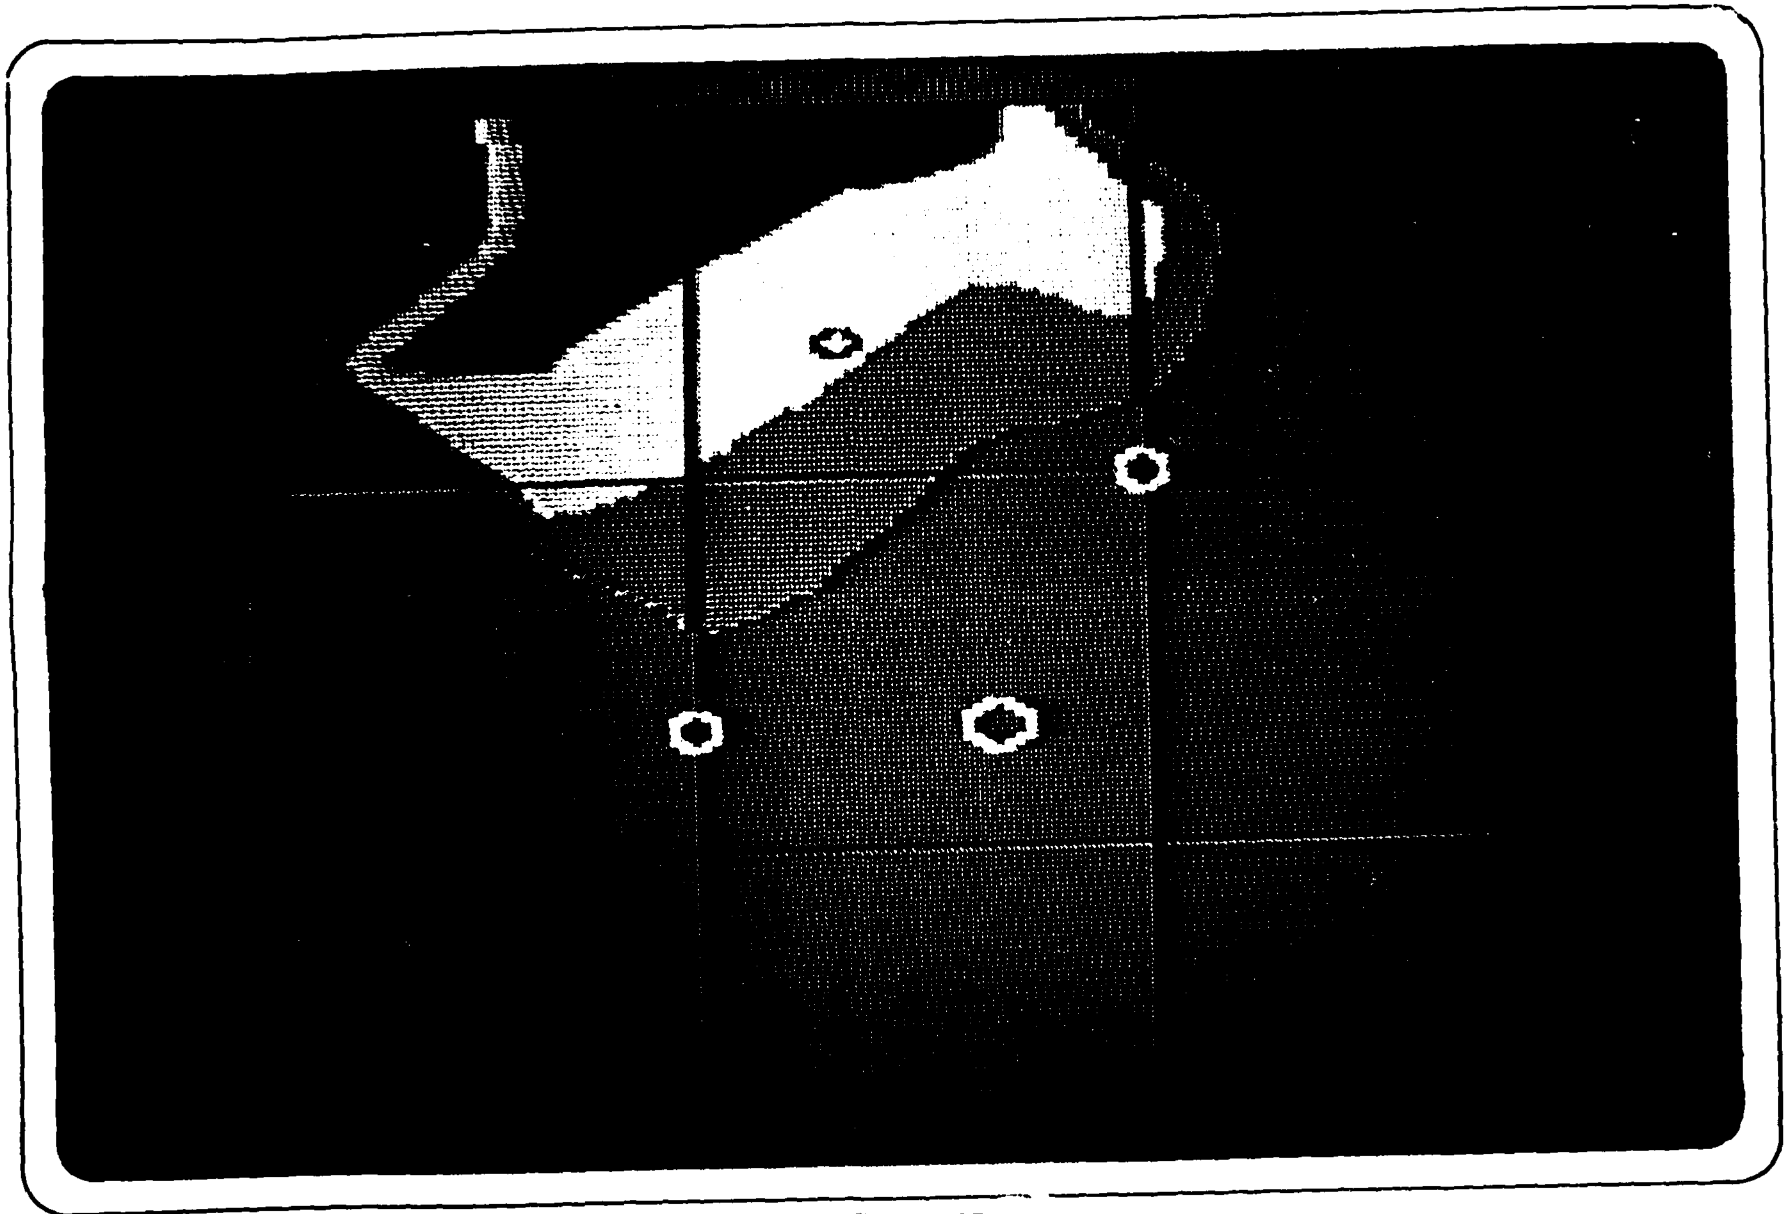

FIGURE 17

LOCATION PROBABILITIES FOR A LOST TRACK THAT REMAINS UNDETECTED FOR 20 MINUTES GIVEN A UNIFORM PRIOR DISTRIBUTION FOR ITS LOCATION

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

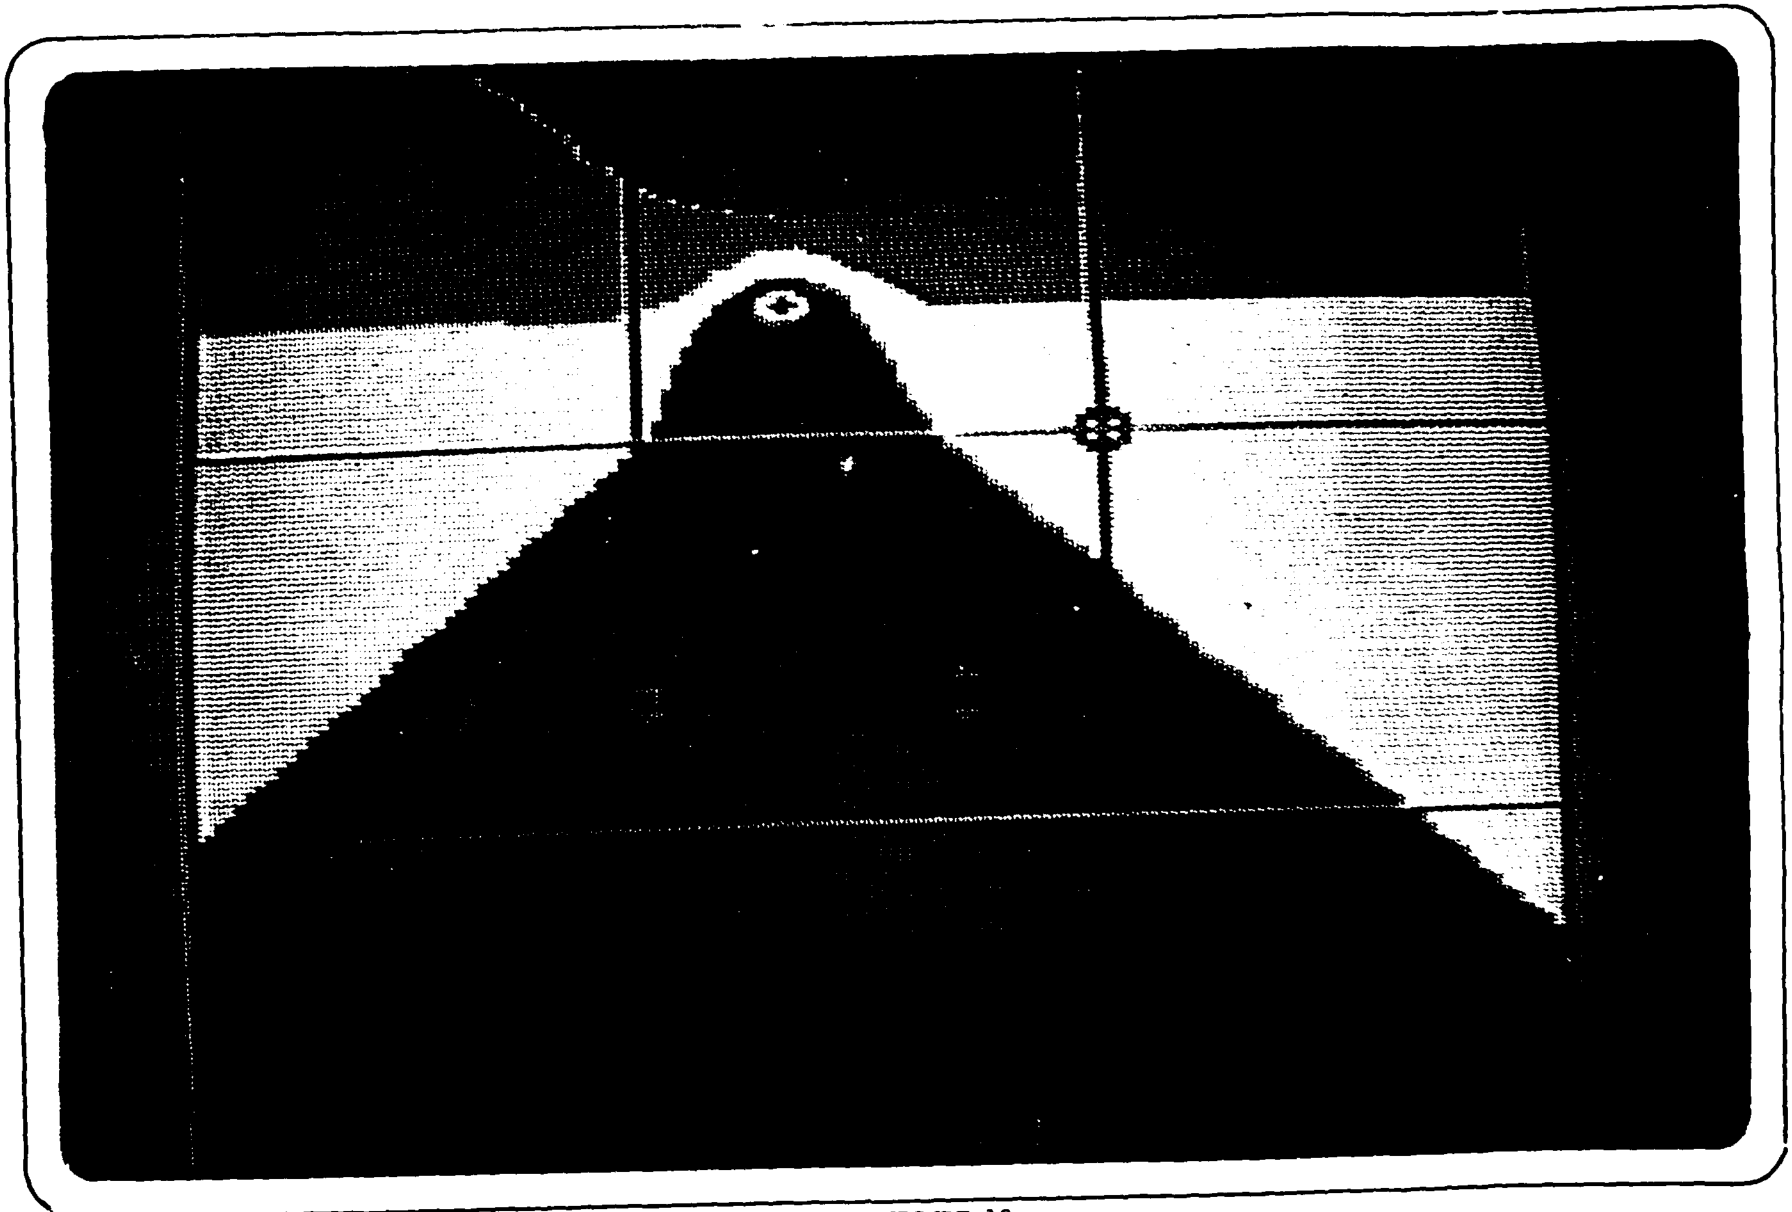

FIGURE 18

A PRIOR DISTRIBUTION OF LOCATION PROBABILITIES FOR A TRACK
THAT WAS TRAVELING SOUTH WHEN IT WAS LOST

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

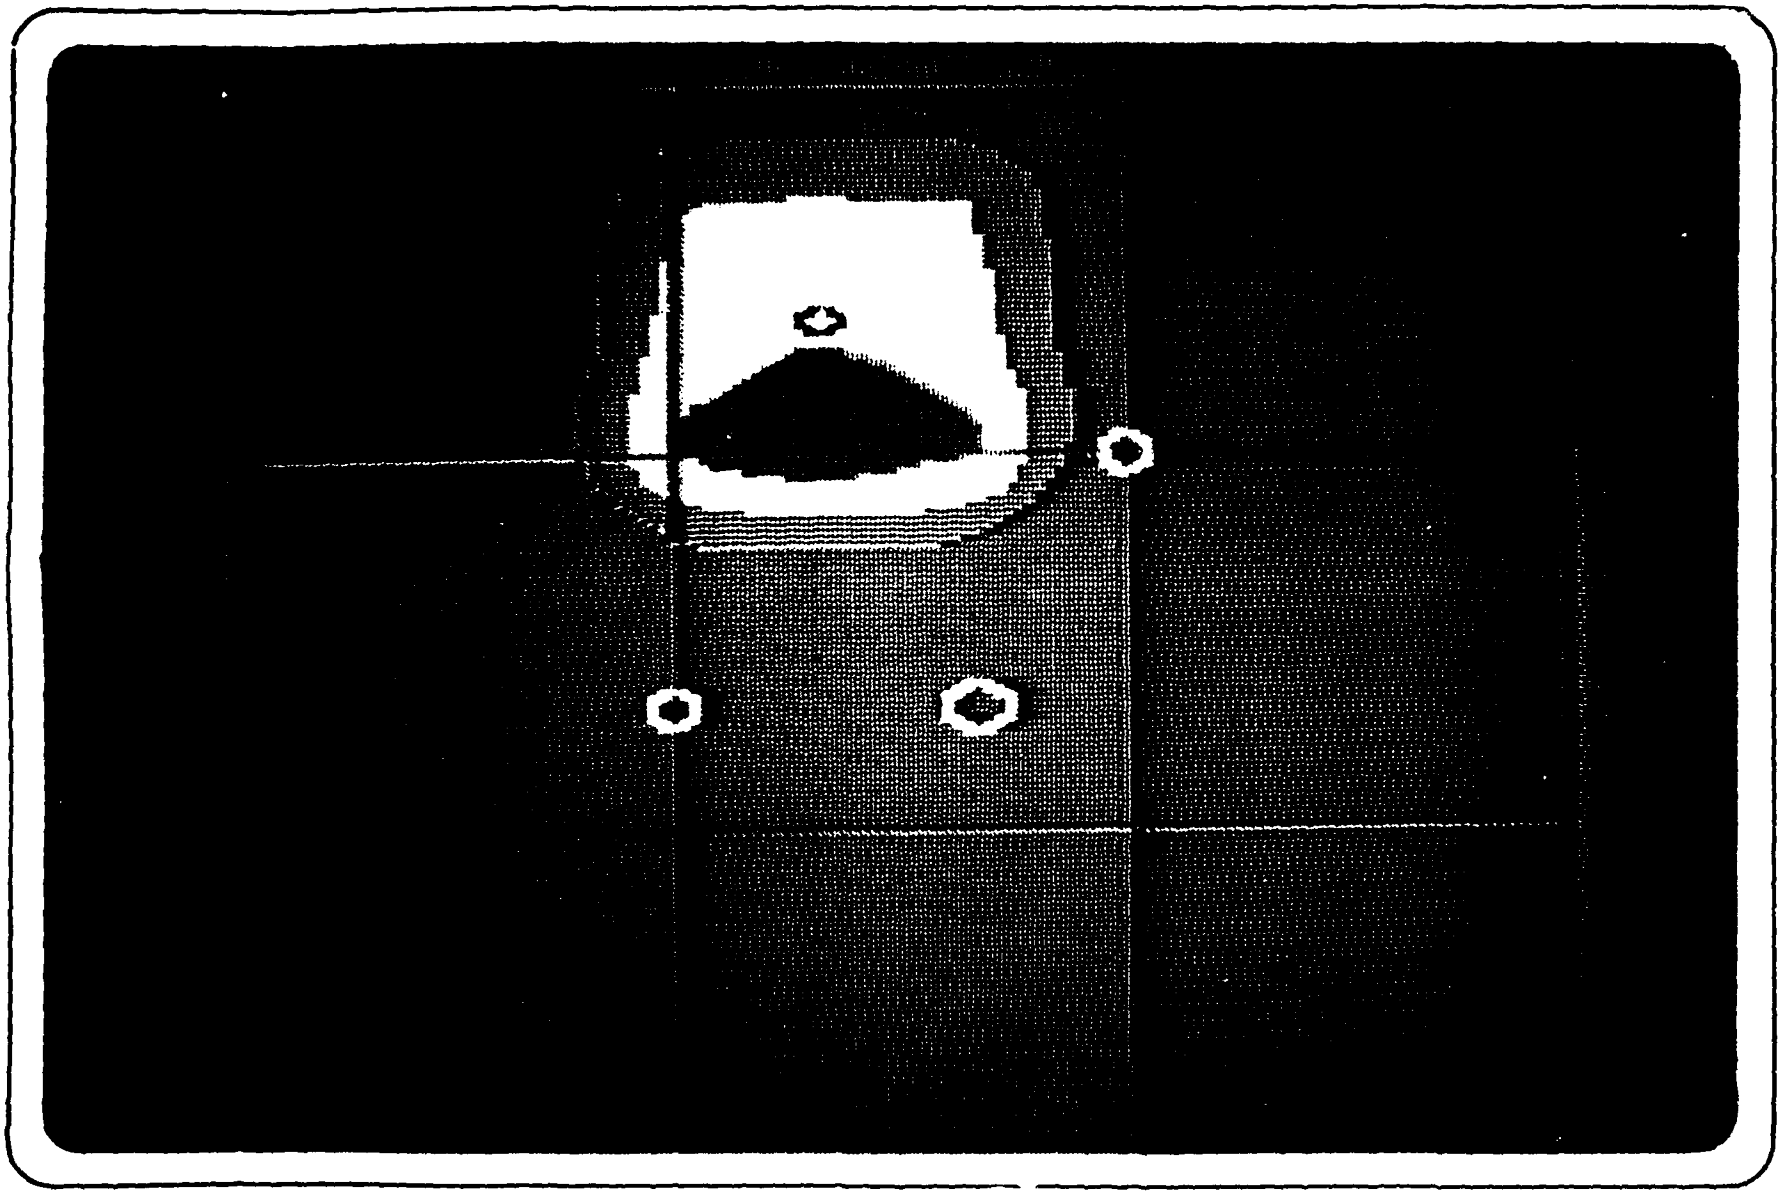

FIGURE 19
LOCATION PROBABILITIES FOR A LOST TRACK THAT REMAINS UNDETECTED FOR 10 MINUTES
GIVEN THE PRIOR DISTRIBUTION FOR LOCATION IN FIGURE 18
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

the prior distribution of Figure 18), but the fact that it has not been detected means that it is unlikely to have traveled as far as possible to the south.

The displays of location probabilities illustrated in Figures 16-19 show the uncertainty associated with a single lost track. However, once the location probabilities have been calculated for several lost tracks (or tracks that have never been detected), it is relatively simple to calculate and display the probability that one or more undetected threats exist at any location. Displays for the location probabilities associated with multiple lost tracks have not been implemented, but they would resemble Figure 19 with several regions colored to show the likelihood of several undetected tracks.

The calculations required to determine the location probabilities of tracks that have never been detected are considerably more difficult than those for a lost track. The difficulty lies in the fact that the number of attacking aircraft or missiles, the times at which they initiate the attack, and the points at which they enter the region shown on the display are not known Probability distributions can be specified for these quantities, and they can be used to generate location probabilities for undetected tracks using the algorithms employed for lost tracks. However, the amount of computation required to do this is significantly greater than that for a lost track. For this reason the implementation of this display has been postponed while other algorithms are investigated. One alternative to the algorithm that generated the location probabilities in Figures 16-19 is based on Markov processes. It is discussed in the Appendix.

First Detection Probabilities

This decision aid displays the probability that a currently undetected track will be detected for the first time at any location near the task force. This information is closely related to the

surveillance penetration and location probabilities. However, first detection probabilities are better suited to the needs of those assigning defensive weapons to threats than penetration or location probabilities because they indicate the areas where defensive weapons may have to be targeted. First detection probabilities also determine the reaction times within which AAW personnel will have to assign and fire defensive weapons. For this reason, first detection probabilities are used by other decision aids to calculate the expected consequences of an engagement.

The probability that a threat will be detected for the first time at a given location is equal to the probability that it can reach that point undetected (i.e., the penetration probability) multiplied by the probability that it will be detected in the vicinity of that point, which is calculated from the detection rate at that location. This information is available from the other decision aids, which means that the display of first detection probabilities requires the same input data and assumptions as the previous aids.

If it is assumed that the enemy will use the minimum detection routes to reach any point near the task force (see Figure 6), and he is equally likely to attack through any point in the vicinity of the task force, the resulting first detection probabilities are shown in Figure 20. This display shows that threats are most likely to be detected for the first time in the annulus around the central ship. Since there is a low probability that a threat can penetrate to the immediate vicinity of the central ship without being detected, the probability of first detection declines in this region.

However, the assumptions used to generate Figure 20 are not very realistic. As discussed earlier, certain attack routes are more consistent with the enemy's tactics, objectives, and knowledge than others. The likelihood that the attack route will pass through any

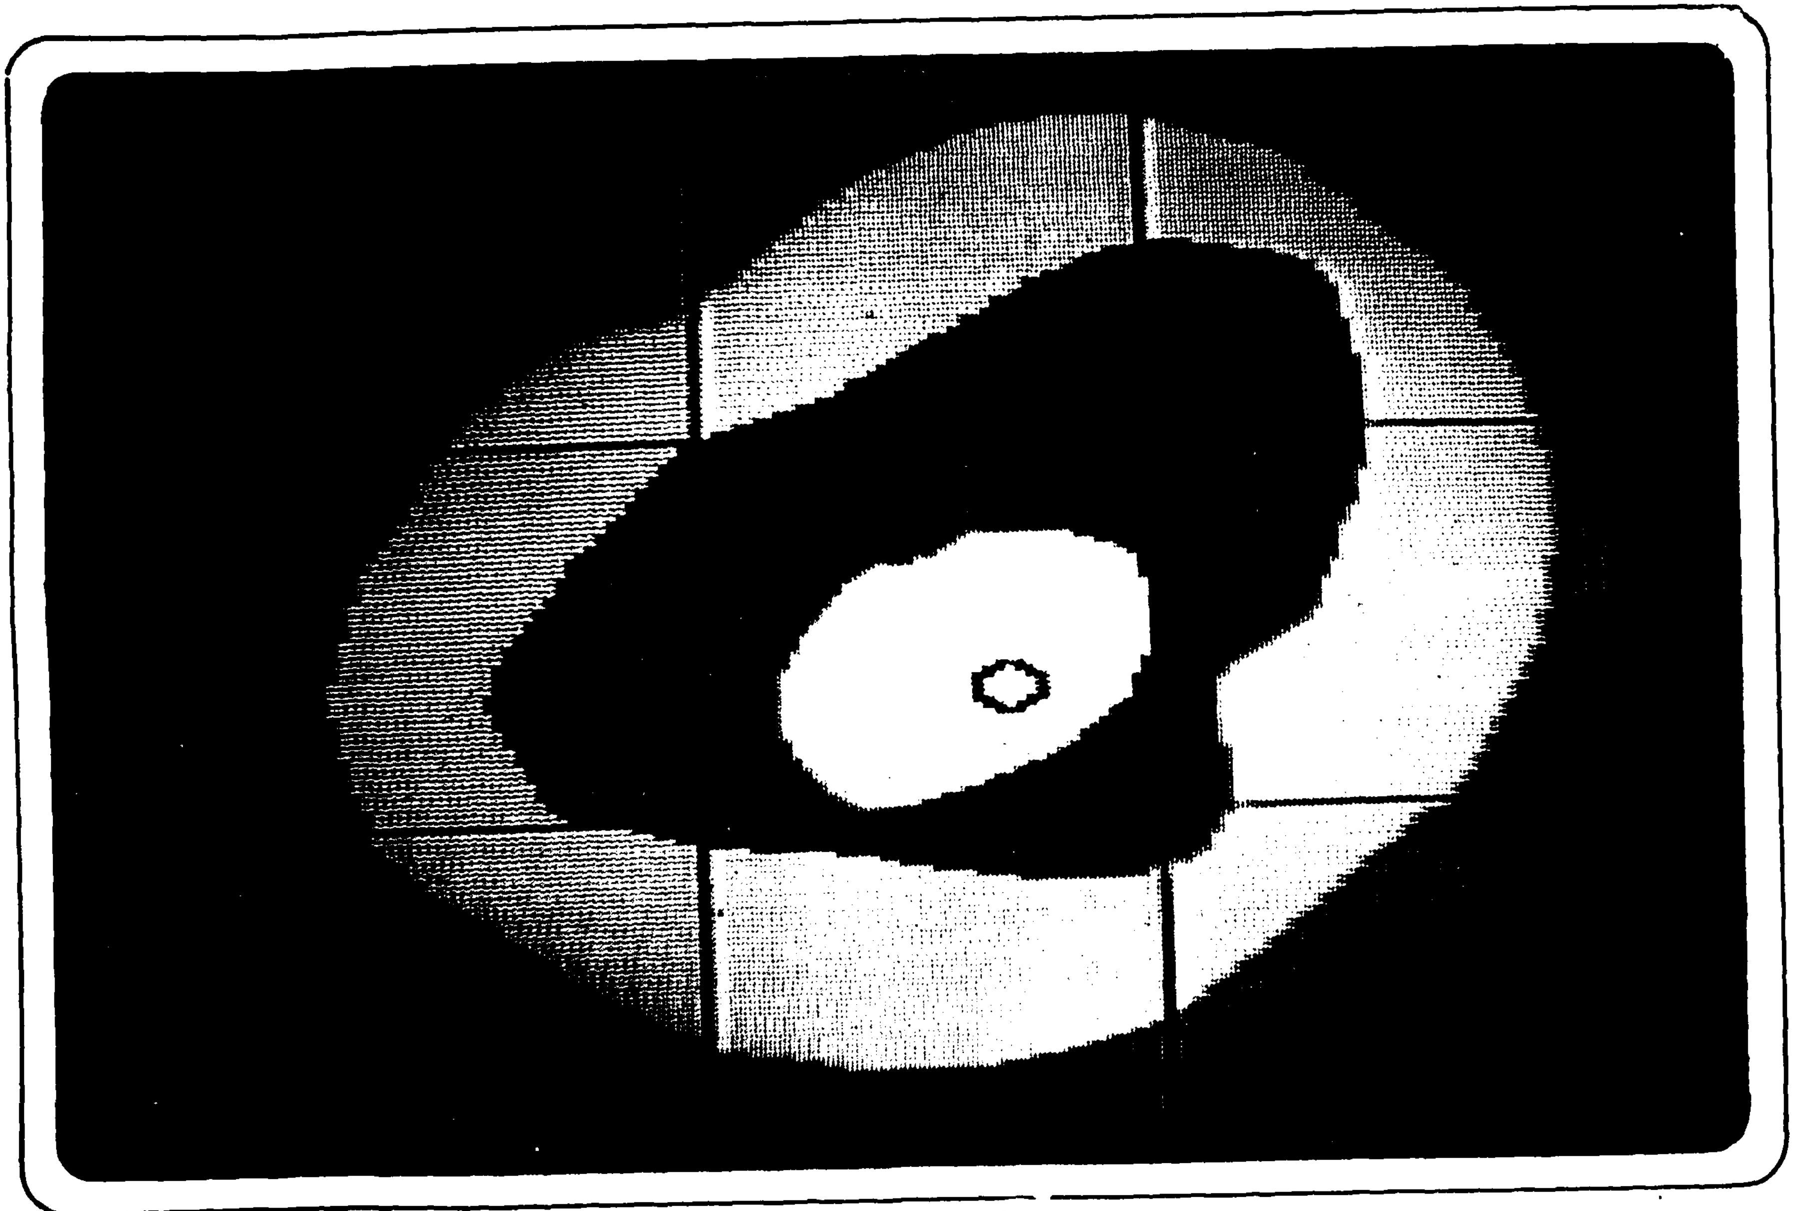

FIGURE 20

FIRST DETECTION PROBABILITIES FOR AN ATTACK FROM ALL DIRECTIONS GIVEN THAT THE ENEMY IS EQUALLY LIKELY TO ATTACK THROUGH ANY POINT NEAR THE TASK FORCE

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

point near the task force is calculated and displayed by one of the aids discussed above (see Figure 9). This information can be used to produce a better estimate of the first detection probabilities. The algorithm used to incorporate this information in the display is described in detail in Section 3 of the Appendix.

Figure 21 shows the display of first detection probabilities if the likelihood of various attack routes is that indicated by Figure 9. In this case the enemy is attempting to reach the central ship while minimizing the probability of detection, he is free to attack from any direction, and he knows the location and performance of the radars. The first detection probabilities in Figure 21 are considerably different than those in Figure 20 because the display in Figure 21 makes use of the fact that the attack is likely to come from the south and east (as shown in Figure 9).

First detection probabilities can also be calculated and displayed for a lost track. However, if the track has been undetected for any length of time, the calculation must be done in two steps. First it is necessary to calculate the probability that the track is currently located at any feasible location (i.e., one it can reach in the time since it was lost) using the algorithms contained in the previous decision aid. Then the aid can calculate the probability that the track will be detected for the first time as it continues from each possible current location. Finally, all of the resulting first detection probabilities can be combined and displayed. This requires considerably more computation than the case of a track that has never been detected, and, as a result, this algorithm has not been implemented. An alternative algorithm based on Markov processes has been explored for this aid as well as the previous one. Its application to first detection probabilities is described in Section 3 of the Appendix.

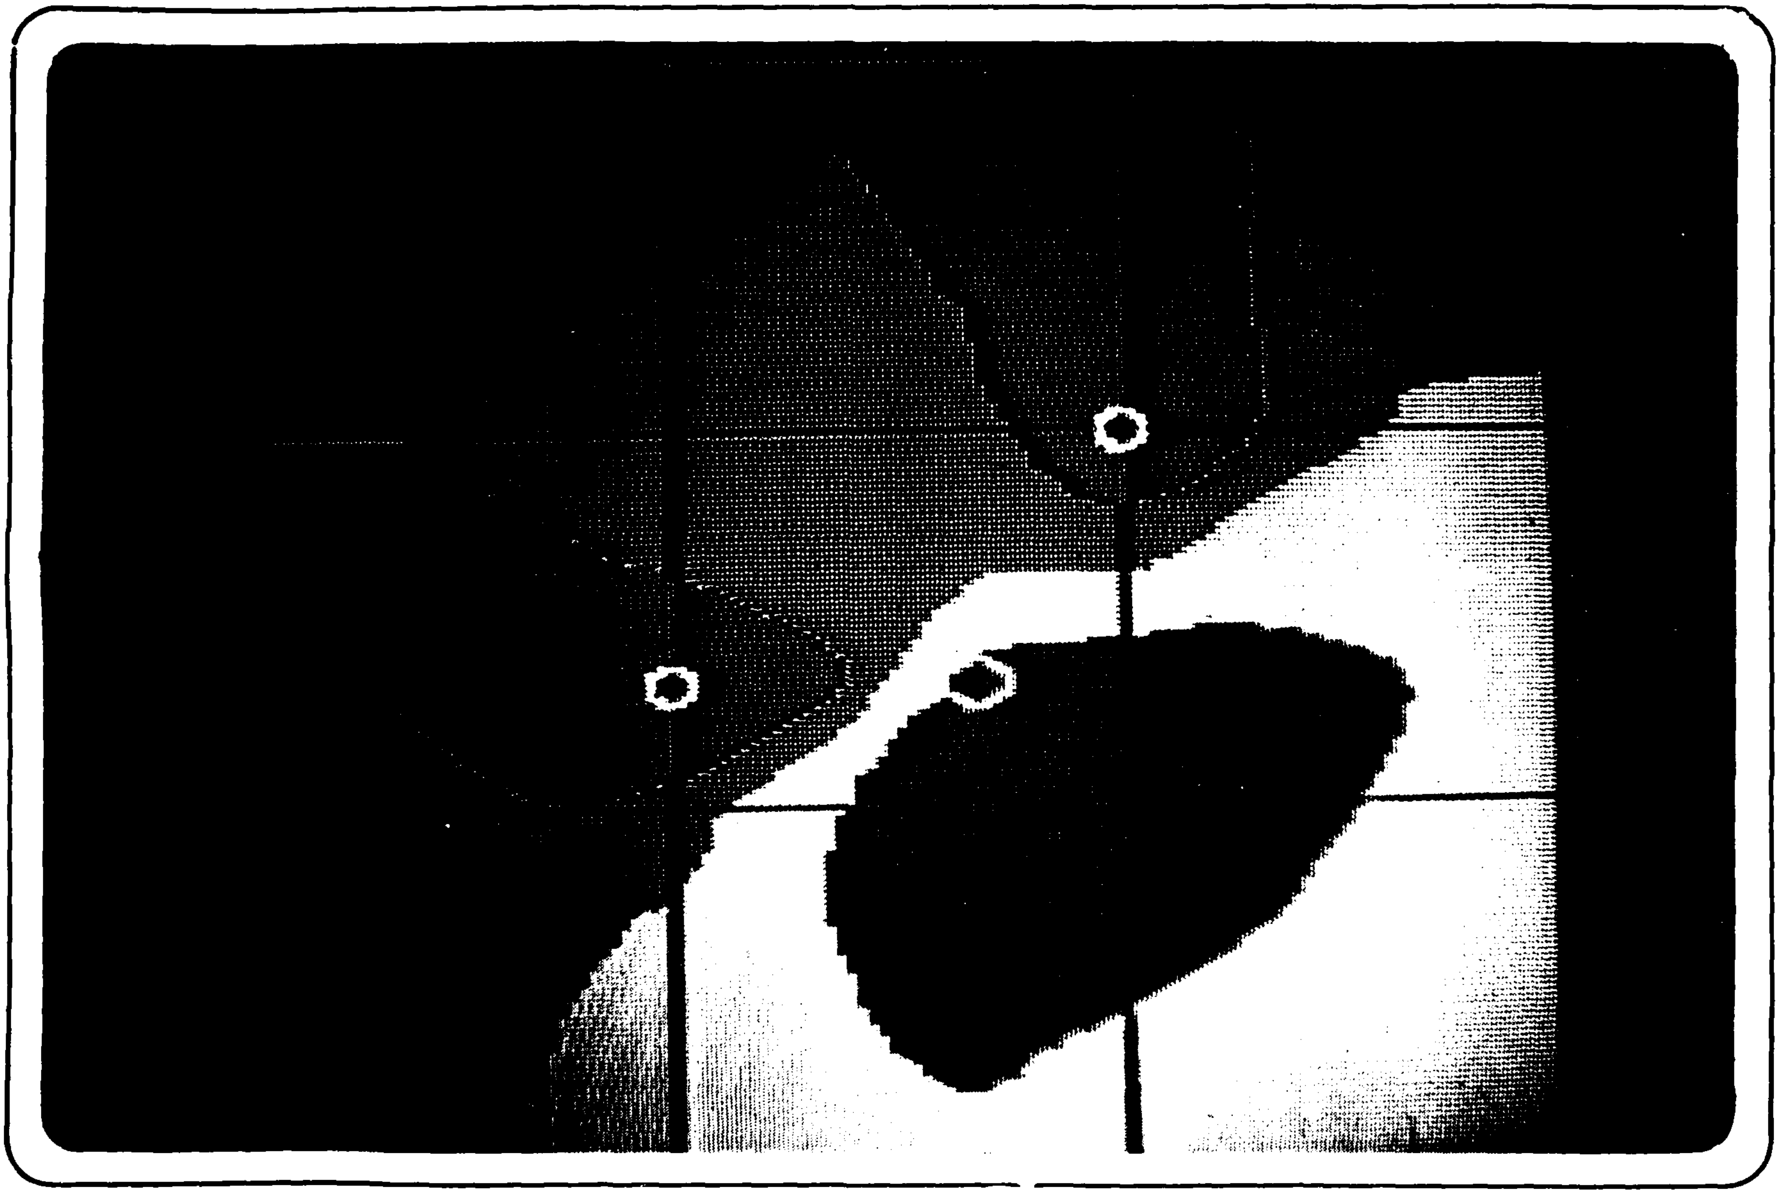

FIGURE 21

FIRST DETECTION PROBABILITIES FOR AN ATTACK FROM ALL DIRECTIONS GIVEN THAT THE LIKELIHOOD OF ATTACK THROUGH ANY POINT NEAR THE TASK FORCE IS THAT SHOWN IN FIGURE 9

Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

Total Penetration Probabilities

This aid combines the first detection probabilities produced by the preceding aid with information about known (i.e., detected) threats and a policy for allocating defensive weapons to determine the probability that a threat can penetrate to any point near the task force without being destroyed. This calculation uses a simple combat model to describe the effectiveness of defensive weapons. It produces a display similar to that for surveillance penetration probabilities (see Figures 11-15), but the display shows the probability of survival rather than nondetection. The inputs, displays, and some of the logic for this aid have been specified, but it has not been implemented.

Like the display of surveillance penetration probabilities, this display is designed to help AAW personnel evaluate the positioning of defensive ships and aircraft. However, total penetration probabilities can be used to judge the effectiveness of the current defensive posture in eliminating future threats, in addition to detecting them.

A typical display of total penetration probabilities is illustrated in Figure 22. It shows the probability that a particular type of threat can reach any point near the task force, including the locations of various ships, given that the enemy can attack from all directions and attempts to avoid detection until a defensive weapon is fired at him. Since this display has not been implemented, the regions of equal penetration probability in Figure 22 may not correspond exactly to the data used to generate the previous displays.

The display in Figure 22 is generated using a simple combat model based on the probability that various defensive weapons can destroy a threat as a function of the distance from the threat to

FIGURE 22
TOTAL PENETRATION PROBABILITIES FOR AN ATTACK FROM ALL DIRECTIONS
Red corresponds to regions of highest probability, and blue corresponds to lowest probability.

weapon at the time it is fired. Typical data is shown in Figure 23, where the probability that a particular defensive missile can destroy an attacking aircraft is plotted for various ranges. This data could be derived from either the results of tests performed with the defensive missile system or a more detailed model of the weapon's characteristics and performance. If detailed combat models are available, they should be summarized by performance curves like that in Figure 23, and not used directly to generate the display in Figure 22. It is more important that the aid produce a reasonable estimate of the penetration probabilities in real time than it is to display the results of a complex combat simulation.

Using data like that in Figure 23 and a weapons allocation policy specified by the user, the aid calculates the number of defensive weapons that can be fired at the threat in the time interval between detection and arrival at the target, the range at which each weapon will be fired, and the resulting probabilities that the threat will reach various points on the way to the target. This calculation is performed for both known and undetected threats.

In the case of undetected threats, the calculation is repeated for each of the locations at which the threat could first appear and weighted by the corresponding first detection probability. Since there may be a large number of points at which a threat could be detected for the first time, this will require more computation than that required for most of the other decision aids. However, several approximation techniques are available to simplify the probability calculations so they can be done in real time. These include proximal analysis and various numerical integration techniques. Further research is needed to determine which solution techniques are best suited to the problem.

This decision aid does not attempt to optimize the allocation of defensive weapons. Instead it shows the effect of a defensive policy established by the user. For instance, the user might specify the

FIGURE 23

EFFECTIVENESS OF DEFENSIVE WEAPONS
AS A FUNCTION OF RANGE

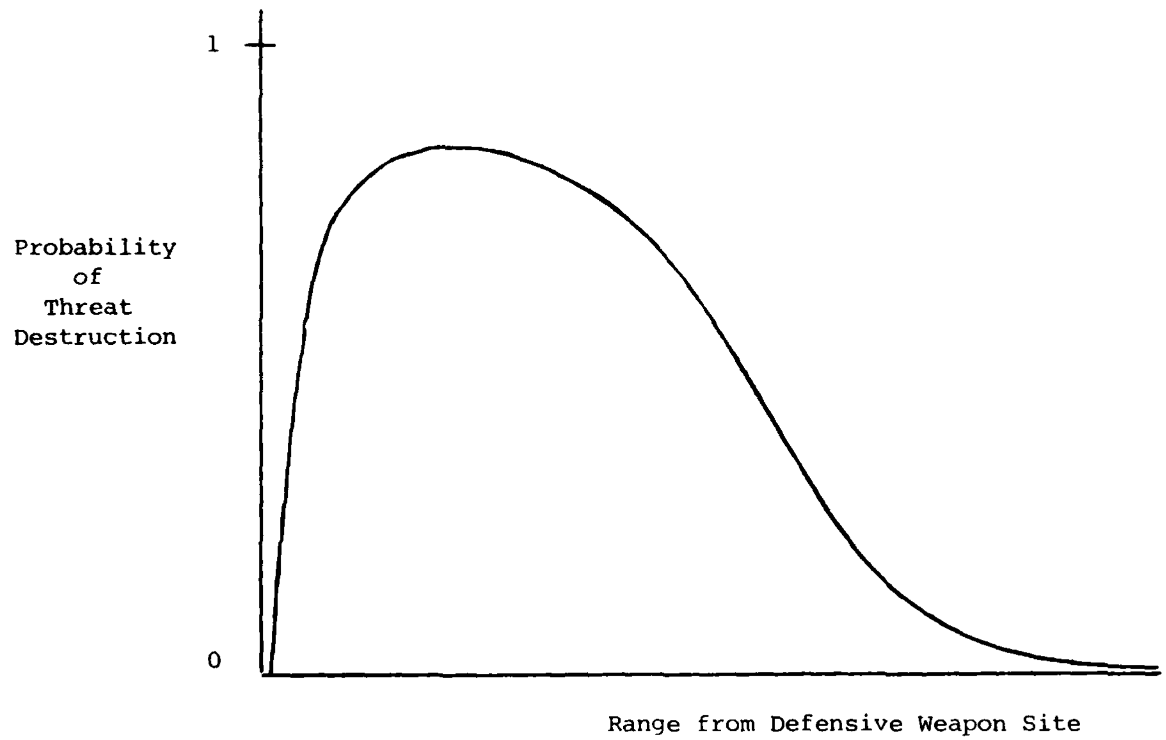

defensive weapons to be used against threats in various areas near the task force, and require that a single defensive weapon be allocated to each threat within a given distance of the task force, followed by a second weapon if the first fails. The decision aid would show the implications of this policy for threat penetration, and, if the results are unacceptable, the user could try a different policy. In this manner, AAW personnel would retain control over all of their resource allocation decisions, and the decision aid would serve to help them test alternative policies and select an acceptable one.

Sources of Expected Damage

This aid uses the same logic and data as the previous one, plus information about the amount of damage that can be caused by various threats, to produce a display of the expected damage (i.e., threat level) associated with both known and undetected threats. The display is generated in real time, and it changes as threats are detected, move to new positions, or encounter defensive weapons. It is designed to support AAW personnel at the unit level who prioritize threats and assign defensive weapons to counter them.

Like the preceding display, this one has not been implemented, but its inputs and displays have been specified. A typical display is shown in Figure 24. It includes both track-specific information (for detected threats) and a threat level defined over the entire area (for undetected threats). Associated with each known (i.e., detected) threat is the expected amount of damage it can cause ships in the task force. In the display the expected damage associated with detected threats is represented by the color of the symbols for the tracks, with red indicating the most dangerous threats. For undetected threats, the expected damage is represented by the colored regions around the task force. According to this figure, the next attack is expected to come from the north.

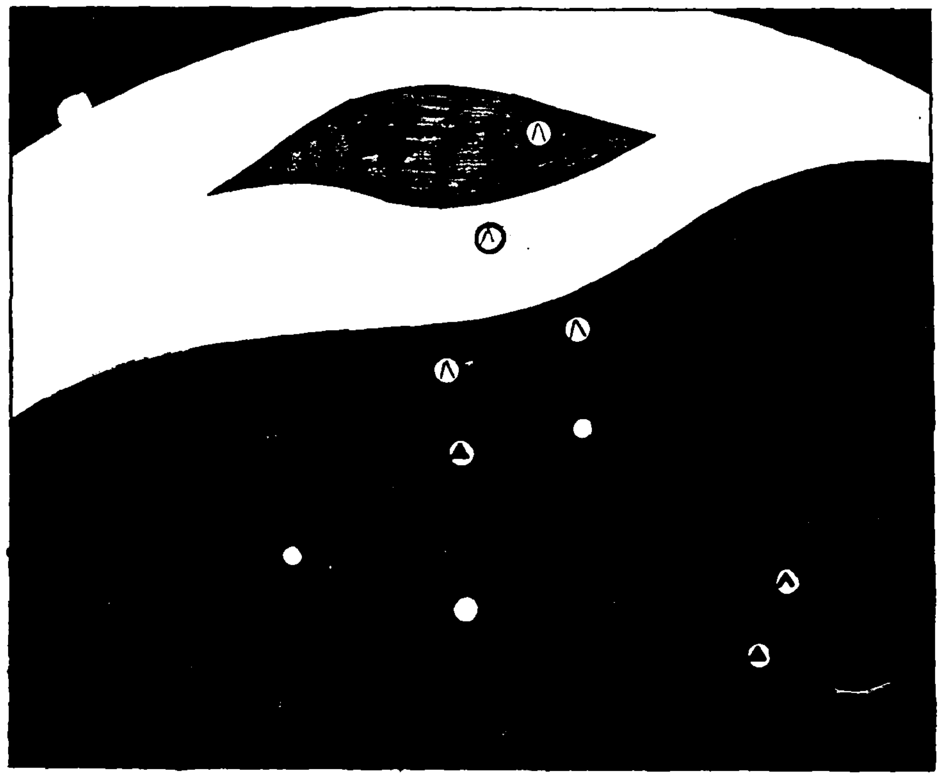

FIGURE 24

SOURCE OF EXPECTED DAMAGE FROM KNOWN THREATS AND AN ANTICIPATED ATTACK FROM THE NORTH

Red corresponds to regions of highest expected damage, and blue corresponds to lowest expected damage. The top of the figure is north.

This display is based on the same combat model used by the preceding aid, and thus requires the same data and assumptions. In addition, it requires a table showing the amount of damage various threats could cause each type of ship in the task force if they can complete the attack sucessfully. This information usually is expressed as the fraction of the ship's combat potential that would be destroyed if the ship sustained a successful attack. This fractional loss is multiplied by the relative value of each ship to produce a measure of damage from the attack.

The preceding decision aid calculated the probability that both known and currently undetected threats would reach specified targets. These probabilities are multiplied by the damage associated with a successful attack to determine the expected damage levels. For known threats, the expected damage levels are translated into colored track symbols. For undetected tracks, the expected damage levels are associated with the locations at which the threats could appear and colored regions are displayed to show locations with similar damage potential. Like the preceding aid, this one is based on first-detection probabilities rather than location probabilities since the allocation of defensive weapons depends on where threats are first detected.

The calculations required to generate this display are similar to those for the preceding aid. In particular, the computations required to generate the expected damage regions for undetected threats will exceed those needed for most of the other aids since the calculation must be repeated for each point at which an undetected threat could appear. However, the approximation technique envisioned for the preceding aid could also be used to generate the display of expected damage in real time.

IV. DIRECTIONS FOR FUTURE RESEARCH

As discussed in the preceding section, a significant amount of work remains to be done in order to complete development of the proposed aids. In particular, implementation of the displays of location probabilities for undetected (anticipated) tracks and first detection probabilities for lost tracks has been deferred because these displays require more computation time than other displays. Alternative algorithms are currently under investigation, including an approach based on Markov processes described in the Appendix.

Two of the decision aids discussed in this report have not been implemented: the displays of total penetration probability and sources of expected damage. However, the current research project has resulted in a specification of the required input data and the form of each display. One goal of future research will be to complete the design and implementation of these displays.

A major requirement for production of these two displays is the development and implementation of a simple combat model. This model will describe the effectiveness of various defensive weapons as a function of the range to an incoming threat and the threat's characteristics. The combat model should include a simple method for specifying a large number of defensive strategies with a few parameters. Navy personnel will select a defensive strategy using these parameters, and the total penetration and expected damage displays will exhibit the implications of the chosen strategy.

With the addition of a combat model, we can revise the objective function used to determine the most likely attack routes. Currently, the relative likelihood of alternative attack routes is based on the enemy's objective of minimizing both the probability of threat detection and distance traveled. However, minimization of the probability of detection is only a surrogate for the true enemy

objective, which is to maximize the damage inflicted on the task force. Detection probability was chosen as a substitute for damage levels until the combat model becomes available because these quantities are correlated: defensive response times are shorter for threats that successfully avoid detection until they are close to their targets. Future research will incorporate the damage potential associated with each path, as specified by a combat model, into the algorithm that predicts the most probable attack routes for incoming threats.

This modification of the choice of routes by incoming threats will affect the appearance, but not the logic or implementation, of most of the experimental displays. The more realistic attack routes will affect the displays of surveillance penetration probabilities, location probabilities, first detection probabilities, total penetration probabilities, and sources of expected damage. (See Figure 3 in Section III).

Continuing development of the aids presented in this report should include the implementation of sensor detection models based on empirical data. In addition, the experimental displays should be tested on a larger computer system, so that algorithmic inefficiencies can be assessed and corrected. One set of unanswered questions deals with the computational savings possible from efficient approximation techniques, particularly for the calculations of first detection and total penetration probabilities.

Displays showing the density of detected, lost, and undetected (anticipated) threats, as well as the density of friendly forces, represent another possible avenue of future research. Such displays could lead in a natural way to status summaries of controlled areas and battle zones. This analysis could be valuable in showing where each side is vulnerable to attack or the repositioning of opposing forces.

APPENDIX: TECHNICAL DESCRIPTION OF ALGORITHMS

1. Detection Rates and Probabilities

For many of the displays described in this report, it is necessary to calculate the probability that an incoming aircraft will be detected as it approaches and passes through the task force area along a specified path. This probability depends on the operating characteristics of the detection devices, environmental factors, and properties of the incoming aircraft.

In this section of the appendix, we derive a formula for the probability of avoiding detection while traveling a short distance in the task force vicinity. Assuming the probability of detection along any portion of a path is conditionally independent (given the sensor detection rate) of the probability of detection elsewhere along the path, nondetection along a path can be approximated by the product of probabilities of traveling undetected over each small segment of the path. In particular, the algorithms described in succeeding sections of this appendix operate under the assumption that realistic flight paths can be approximated by piecewise linear paths over a sufficiently fine grid of points.

Inputs

- Rate of detection of the approaching aircraft or missile.
- Velocity of the incoming aircraft or missile.

Output

- Probability of traveling in a straight line between two neighboring points, without being detected.

Assumptions

- The detection rate is sufficiently smooth that paths over a short distance can be approximated by a straight line.
- The aircraft or missile travels at a constant velocity (i.e., its maximum sustainable velocity).
- Given all physical data and detection characteristics, detection attempts ("looks") by a configuration of sensors are independent of previous looks.

Data Sources

The logic required to calculate the detection rates from physical data has not been developed as part of this project, and a separate research effort will be needed to specify appropriate algorithms. In principal, the detection rate can be calculated by combining the following information for each sensor:

- Type of device (radar, IFF, etc.);
- Detection characteristics (signal/noise ratio, directional biases, etc.);
- Status (ON/OFF, passive/emitting, reliability, damage sustained, etc.);
- Location;
- Environmental conditions (weather, chaff, jamming, etc.), and their effect on the sensor.

All of this information is either currently available on the Navy Tactical Data System (NTDS), or will be available in future systems. A simplifying assumption that could be used to calculate the overall detection field from this data is that detection by one device is conditionally independent (given environmental conditions and status) of detection by other sensors. However, it is not necessary to make

this assumption, nor do any of the algorithms or aids require this conditional independence. A more practical model would estimate the detection rate for a group of sensors based on empirical data.

In addition to characteristics of the sensors, the detection rate may be influenced by properties of the threat. For instance, certain aircraft (e.g., stealth aircraft) are less likely to be detected by radar devices. An estimate of the effects that the design of the threat has on the detection rate should be made based on available empirical evidence.

In the absence of sufficient information to specify the exact detection characteristics of each sensor device, a convenient functional form can be substituted for the detection rate. The correct functional form can be determined from the large quantities of empirical data available for all Navy detection devices. For testing the experimental displays, ADA chose to use a detection rate that is inversely proportional to the square of the distance from the device. In this formulation, the capabilities of each device are completely specified by its location, a parameter denoting its strength at the source, and a second parameter representing the dispersion of its detection potential over distance. The detection rates derived from this model, although they are not based on empirical data, nonetheless produce plausible displays.

Formulation

We wish to determine the probability of traveling along a specified path between the two points z_0 and z_1, without being detected. Let $z(t)$ represent the position of the threat at time t, with $z(t_0) = z_0$ and $z(t_1) = z_1$. Then the detection rate r at any point along the path is defined by:

$$r(z(t))dt = \text{Prob.}\{\text{detection occurs in } (t, t+dt) \mid \text{no detection prior to time } t \text{ and location } z(t) \text{ at time } t\}.$$

The assumption implicit in this definition is that the detection rate at $z(t)$ does not depend on the path taken to reach $z(t)$. In particular, the success of a detection attempt ("look") at time t is independent of the number, location, and timing of all previous looks.

Let $P(t)$ be the probability of remaining undetected through time t. The following derivation produces an expression for $P(t)$.

$$P(t + dt) = P(t)\ (1 - r(z(t))dt)$$

$$\frac{dP(t)}{dt} = -r(z(t))P(t)$$

$$P(t) = \exp\left\{- \int_{t_0}^{t_1} r(z(t))dt\right\}$$

Next, represent the path as a function on the two-dimensional plane. That is, let $z(t) = (x(t), y(t))$. Assuming monotonicity of the path in the first dimension (since we will later assume that the path can be approximated by a straight line, there is no loss of generality in this assumption), a simple change of variables yields:

$$P(t) = \exp\left\{- \int_{x_0}^{x_1} \frac{r(x, y(x))}{\frac{dx}{dt}} dx\right\}$$

At this point, we introduce the assumption of constant line velocity V, where $V = (\dot{y}^2 + \dot{x}^2)^{\frac{1}{2}}$. With this assumption we can derive an expression for $\frac{dx}{dt}$ and use it in the formula above:

(1) $$P(t) = \exp\left\{- \int_{x_0}^{x_1} \frac{r(x, y(x))}{V} \left(\frac{dy}{dx}^2 + 1\right)^{\frac{1}{2}} dx\right\}$$

The difficulty of evaluating this expression depends on the functional forms of the detection rate r and the path $y(x)$. For simplicity, we assume that the path can be approximated by a straight line. This assumption is quite reasonable in the context of our problem, since actual flight paths tend to be approximately linear over short distances. Under this assumption, the path derivative is a constant. Then for simple functional forms of r, (e.g., detection rate for each sensor inversely proportional to the square of the distance), the integral in Equation (1) can be evaluated directly. In more complex cases, numerical integration techniques must be used.

2. Penetration Probabilities and the Dynamic Programming Algorithm

The dynamic programming (DP) algorithm is used to find an optimal path to each point in the vicinity of a task force from a fixed set of starting points. If the set of starting points is outside (or at the edge of) the detection range, then the probability of flying along the optimal path undetected is the penetration probability. If a single location is chosen as the starting point, the optimal paths produced by DP can be used to show penetration probabilities for a lost track starting at that point. Another application of the DP algorithm is to discover all optimal paths ending at a particular target point. This information is used to determine the relative suboptimality of various paths to a target (see the next section of this appendix).

There are actually two similar problems that can be solved using the DP approach. One problem is to determine the minimal detection paths, the other is to find the paths which optimize an objective function containing terms for minimizing both detection probability and flight time. For ease of discussion, the minimal detection case will be described first, although the other case is a generalization that subsumes it. A second generalization of the problem, to the case where enemy information about the detection field is imperfect, is developed at the end of this section.

Problem Statement

Sensors are attempting to detect approaching aircraft. At each point x, there is a detection rate $r(x)$ defined by the current operating characteristics of these sensors and environmental conditions (weather, chaff, jamming, etc.). Given an initial set of probabilities for reaching each point undetected (zero for points not

in the starting set), what is the maximum probability of traveling undetected to each point? Which path results in this minimum probability of detection?

Inputs

- A starting set of points with associated (non-zero) probability that they can be reached undetected. (For many applications, this probability is 1.)
- Rate of detection of approaching aircraft.
- Maximum sustainable velocity for incoming craft.

Outputs

- Optimum paths from the starting set to each point in the task force vicinity.
- Probability of reaching each point, along an optimum path, without being detected.

Assumptions

- The detection rate is sufficiently smooth that the minimum detection path between any two points can be approximated by a piecewise linear path.
- Approaching aircraft fly at a constant velocity (typically their maximum sustainable speed in an attack situation).

Data Sources

The detection rate can be calculated from information about detection characteristics, location, current operating status, and environmental conditions for each sensor, as described in the first section of this appendix. All of this information is either currently available on the Navy Tactical Data System (NTDS), or will

be available in future systems. The assumption typically used to calculate the overall detection field from this data is that detection by one device is conditionally independent (given environmental conditions and status) of detection by other sensors.

The maximum sustainable velocity for incoming craft can be estimated by Navy personnel from intelligence information.

Formulation

Lay down a rectangular grid of points covering the entire task force vicinity or detection area. Flight along any path will be approximated by straight line travel between neighboring points of the grid. Each point has eight neighbors as depicted in Figure A.1, and the entire grid has N points.

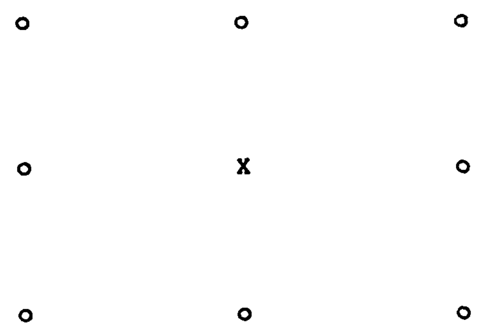

Figure A.1: The 8 Neighbors of Any Point X Located on a Rectangular Grid

Let $p(i,j)$ be the probability of traveling undetected along the straight line from point i to neighboring point j. This quantity is calculated directly from the sensor detection rate and velocity of approaching aircraft (see the first section of this

appendix). Let $V(i)$ represent the probability of reaching point i, without being detected, along the best path found so far. The initial value $V_0(i)$ for this function is normally zero if i is not in the starting set; otherwise it is a value specified by the inputs.

The DP Algorithm

1. Initialize $V(i) = V_0(i)$, for all points $i = 1,2,\ldots,N$. Set FLAG = TRUE. Set $B(i) = i$, for all i.

2. Set $V(i) = \text{Max}\{V(i),\ p(j,i)\cdot V(j) \mid j \text{ is a neighbor of } i\}$ for $i = 1,2,\ldots,N$. If this step changes $V(i)$, set FLAG to FALSE and set $B(i)$ to the index of the neighbor j which optimizes the above expression.

3. If FLAG = TRUE, then halt, otherwise continue to Step 4.

4. Set FLAG = TRUE and go to Step 2.

On every iteration, this algorithm determines a path to each point i and evaluates the probability of traveling undetected along the path. Paths are found by starting at i and finding the best path to a point in the starting set that goes through some neighbor of i. In particular, the algorithm uses the following recursive equation:

$$V^*(i) = \max\{V_0(i),\ p(j,i)\cdot V^*(j) \mid j \text{ is a neighbor of } i\},$$

where $V^*(i)$ denotes the maximum penetration probability for any path from the starting set to point i. If the optimal path to each neighbor is known, then the best path to i must be discovered on the next iteration. The best route from each neighbor is determined simultaneously by the algorithm. On each individual iteration an approximation (the best path found so far) is used so that the entire path to point i need not be examined.

After termination, the optimal path to i can be traced backwards, from i to $B(i)$ to $B(B(i))$, etc., until a point s is reached with $B(s) = s$. This endpoint must be in the starting set.

The DP algorithm converges monotonically to the maximum penetration probabilities, $V^*(i)$. In fact after k iterations, every path from the starting set to point i consisting of at most k straight lines between neighbors will have been implicitly considered. Since the longest optimal path can pass through at most N points, the algorithm will terminate after at most $N+1$ iterations.

The termination criterion of Step 3 states that the optimum values have been found if they do not change from one iteration to the next. This is a valid termination criterion because if a better path to i exists, one can prove that a better path to some point on that path must be found in Step 2.

Implementation

The application of this algorithm to the stated problem can be very efficient. Minimum detection paths can be expected to go from any starting point to an interior point by avoiding areas with a high density of sensors. Thus ordering schemes similar to that shown in Figure A.2 are used to improve the speed at which paths to interior points are found in Step 2. In practice, the DP algorithm normally terminates with optimal paths after only 2 or 3 iterations.

In order to calculate $p(i,j)$, the probability of traveling in a straight line from i to neighbor j without being detected, we must assume that the incoming aircraft travels at a constant line velocity, V. If V is the maximum sustainable velocity for the plane, then $p(i,j)$ will be the maximum probability of escaping detection along the straight line path. The procedure for

AD-A104 369 APPLIED DECISION ANALYSIS INC MENLO PARK CA F/G 15/3
DECISION AIDS FOR NAVY TACTICAL ANTI-AIR WARFARE.(U)
JUN 81 A C MILLER, L J POLLENZ, F SHIH N00014-80-C-0405

UNCLASSIFIED NL

2 of 2
AD
A104369

END
DATE
FILMED
0 81
DTIC

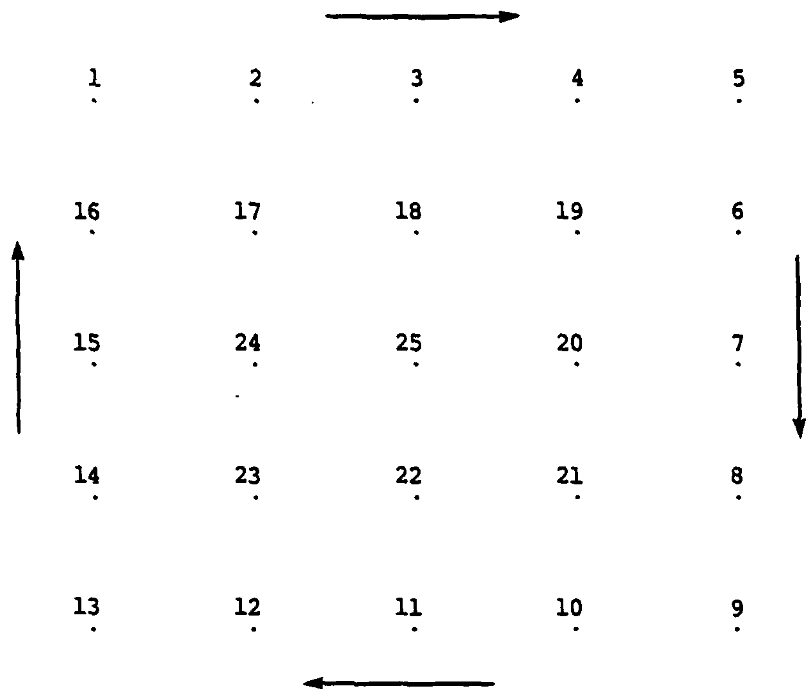

FIGURE A.2: ORDERING SCHEME FOR PATH FINDING ALGORITHM

calculating p(i,j) from physical data is described in the section on detection rates and probabilities at the beginning of this appendix.

Simple Example (Minimal Detection Problem)

Consider a grid consisting of nine points, as in Figure A.3, with a single radar located at the center of the grid. (In general, sensors may also be placed anywhere between grid points.) Let the starting set be the two points labeled 1 and 8 in this figure, with initial values of one. Assume that the radar field is symmetric about the center of the grid, and that the values for traveling in a straight line are given by:

$$p(1,2) = .8,\quad p(2,8) = .5,\quad p(2,9) = .4,\quad p(1,9) = .25.$$

The detection probabilities for all other pairs of points are determined by symmetry. These values correspond to a weak radar under the assumption that the radar detection rate falls off with the square of the distance from the sensor.

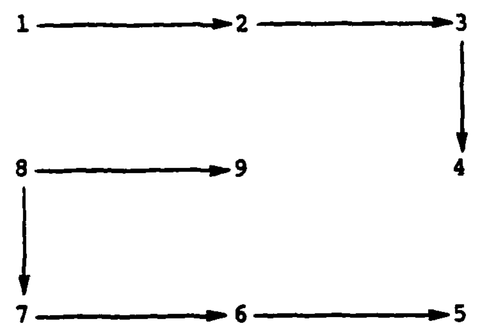

Figure A.3: A Simple Example of Minimum Detection Paths Found by the Dynamic Programming Algorithm

The first iteration of the DP algorithm produces the following paths from the seven points not in the starting set (* denotes optimality):

Endpoint	Value	Path
* 2.	.8	(2,1)
* 3.	.64	(3,2) then path from 2
* 4.	.51	(4,3) then path from 3
5.	.41	(5,4) then path from 4
6.	.5	(6,8)
* 7.	.8	(7,8)
* 9.	.4	(9,8)

On the second iteration, the direction of circulation shown in Figure A.2 is reversed (this alternation of ordering occurs throughout the algorithm, and considerably speeds execution). The path from point 6 starts with (6,7), and the path from point 5 starts with (5,6). These paths are optimum, with values .64 and .51 respectively. Thus the third iteration of the algorithm will find no improvements and termination will occur. The optimum paths are displayed in Figure A.3. Note that these routes shun the central region as much as possible, and therefore do not minimize the total distance traveled.

Generalized Objective Value

As discussed in Section III, it more realistic to assume that the enemy chooses his flight paths based on objectives other than minimizing detection probabilities. Since his ultimate goal is to inflict damage on the fleet, he may wish to minimize the defensive reaction time for Navy vessels. This involves both avoiding detection and decreasing flight time.

Both of these objectives can be incorporated into the DP algorithm by changing the interpretation of $V(i)$. Define $V(i)$ to

be the objective value along the best path to point i found so far. In particular, let

$$V(i) = c^{(m_1 + \sqrt{2}m_2)}\ P(i),$$

where $P(i)$ = probability of reaching i undetected along the chosen path,

m_1 = number of horizontal or vertical line segments along this path,

m_2 = number of diagonal line segments along the path, and

c = trade-off value relating each additional horizontal path step to a percentage increase in detection probability.

The DP algorithm is changed only in Step 2; the optimization in this step becomes:

$$V(i) = \text{Max}\ \{V(i),\ c(j) \cdot p(j,i) \cdot V(j)\ |\ j \text{ is a neighbor of } i\},$$

where $c(j) = c^{\sqrt{2}}$ if j is a diagonal neighbor of i, and $c(j) = c$ otherwise.

The relative importance (to incoming threats) of avoiding detection versus minimizing flight time is an additional input required for this generalization of the DP model. It is expressed as the value of the trade-off parameter c. As c approaches zero, distance becomes all-important, and the shortest, straightest paths will be optimal. At $c = 1$, the minimal detection paths are found, exactly as before. For intermediate values of c, paths will be chosen that are moderately straight and yet avoid the heaviest areas of sensor concentration. See Figure A.4 for an illustration of the effect of changing this model parameter. The source of this information would be intelligence estimates based on the enemy's prior tactics.

FIGURE A.4 INFLUENCE OF THE PARAMETER ON FLIGHT PATH

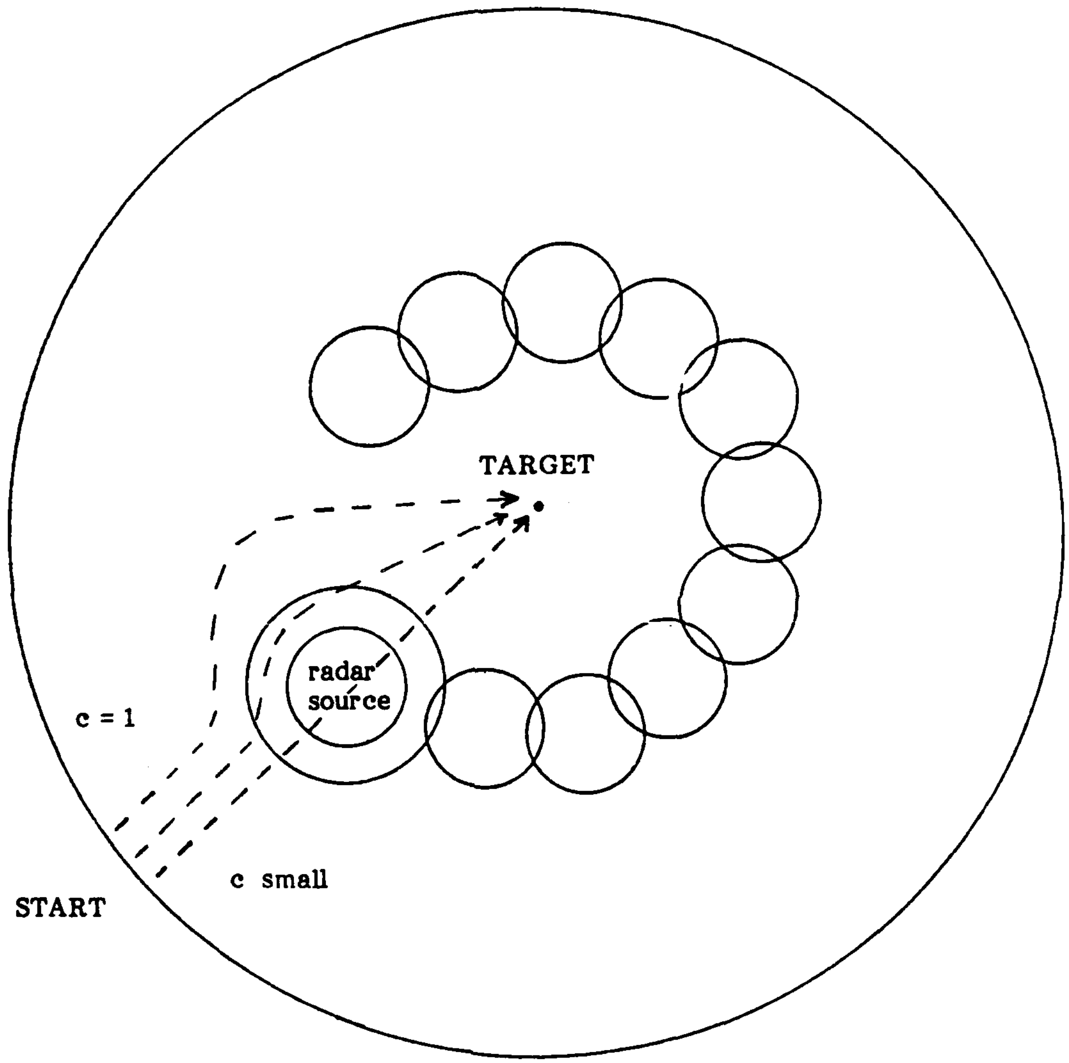

The smaller the value of the parameter c, the more direct the flight path will be.

<u>Example</u>

Suppose the trade-off value is c = .5 for the nine-point sample problem given earlier in this section. This implies that minimizing flight time (flying straight paths) is quite important to the attacking aircraft. The optimum paths in this case are shown in Figure A.5 below. Note the shorter routes taken to reach points 4, 5, and 6.

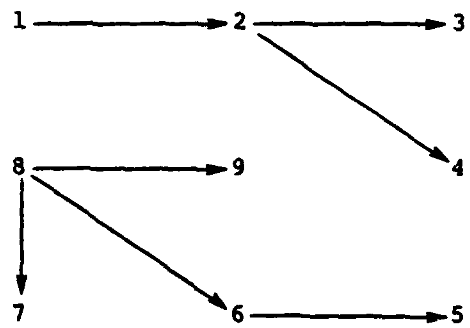

Figure A.5: A Simple Example of Optimal Paths Found by the Dynamic Programming Algorithms When the Enemy's Objective is to Minimize Detection and Distance Traveled

The effect of introducing a value for the path length into the objective function is also illustrated by the comparison of Figures 6 and 7 in Section III of the main report.

Generalization to Perceived Rates of Detection

The enemy may not be certain of the exact strength and placement of Navy sensors. The less sophisticated the enemy's equipment for discovering the whereabouts of emitting devices, the more imperfect or "fuzzy" his perception of the detection rate will be. In addition, he may be unaware of passive detection devices altogether, and thus will not attempt to avoid them.

In order to account for these two effects, the dynamic programming algorithm should force incoming planes to choose their paths based on the perceived detection rate. However, the penetration probability for the chosen paths still is evaluated using the true detection field.

To determine the perceived detection rate $r'(x)$, probability distributions are used to describe intelligence estimates of the enemy's state of information about the location of each of the emitting sensors. The mean of each distribution would be the actual sensor location. In most cases the location distribution would be symmetric about its mean, but directional biases can be incorporated if evidence supports their existence. Then $r'(x)$ is the expected value of the detection rate at x given these probability distributions.

Since the functional form for the sensor detection rate is nonlinear, the "fuzzy" detection rate $r'(x)$ will not be identical to the original detection rate $r(x)$. There will be corresponding changes induced in the probability of traveling undetected between neighbors since this quantity is computed directly from the detection rate (see the first section of this appendix for details). This may cause the dynamic programming algorithm to choose different paths for the perceived detection rate than for the true detection field.

Implementation of Perceived Detection Fields

Uncertainty about the location of an emitting radar for a fixed observer position is equivalent to uncertainty about the observer location for a fixed radar position. This is true because the detection rate depends only on the relative positions of sensors and threats, and environmental conditions between each detection device and an incoming object. Therefore, the perceived detection rate $r'(x)$ can be determined by calculating the expected value of the true detection rate (compensated for unanticipated passive detection) using the probability distribution for radar location over an area centered at x. The size of the area receiving a non-negligible weight corresponds to the spread or variance of the distribution function. The larger the area, the more uncertainty exists, and the fuzzier the observer's picture of the detection field will be.

A good approximation to $r'(x)$ can be obtained by taking an appropriately weighted average of the true detection rates at a discrete set of points within a given radius of x. In particular, a weighted average of the field at a point and its neighbors provides one representation of the perceived detection field.

If the observer's uncertainty about the detection rate is the same for all sensor devices, then the effect on the probability of traveling undetected between neighboring points is relatively simple. In this case, the transformation is:

$$p'(i,l_i) = p(i,l_i)^{(w_i)} \cdot \prod_j p(j,l_j)^{(w_j)}$$

where l_i is a neighbor of point i in a specified direction, and w_j is the nonnegative weight assigned to neighbor j of i, such that $w_i + \sum_j w_j = 1$.

The choice of weights (or equivalently, the probability distribution over sensor location) depends on the quality of information available to attacking aircraft, as determined by the sophistication of their devices for detecting and locating emissions. A single input parameter is used to describe this level of sophistication in the current implementation of the aid. The parameter corresponds to the variance of a symmetric distribution about the actual sensor location. The source of this parameter would be a subjective assessment by Navy personnel, based on intelligence information.

One possible (although conservative) setting for this parameter is to assume the enemy has perfect information; this is equivalent to the original version of the DP algorithm. As the quality of information is reduced by lowering the sophistication parameter, the perceived location of the sensors and the detection rate become fuzzier. Thus circuitous routes for incoming threats trying to avoid detection become less likely, particularly if an objective value placing some importance on flight time is included.

3. First Detection Probabilities and the Relative Suboptimality of Paths

For the display of first detection probabilities, the probability that each path will be chosen by incoming aircraft must be evaluated. This calculation requires some of the same assumptions concerning enemy behavior as the dynamic programming algorithm, but it is recognized that suboptimal (especially near-optimal) routes may be chosen. Reasons for choosing a suboptimal path are:

1. Imperfect information (both our intelligence information and the enemy's);

2. Imperfect estimation of the trade-off between detection avoidance and flight time;

3. Other objectives, such as avoiding defensive strength or staying in formation;

4. The use of a mixed strategy randomly selected from a set of good paths in order to make enemy behavior less predictable.

The problem is to find the relative likelihood of the various paths to a specific target, and then aggregate over all possible targets. We have developed two different approaches for solving this problem. One method compares the objective value obtained along the best path to the target which passes through a particular point with the objective value for the optimal, unrestricted route to the target. The objective value used here will be the same as for the generalized dynamic programming algorithm, including a term for minimizing distance along the path and the possibility of imperfect enemy information about sensor locations. The second approach uses an assumption of localized random behavior to allow suboptimal paths to be chosen. This probabilistic solution technique uses both dynamic programming and Markov chain theory.

Inputs

- Sensor detection rate.
- Maximum sustainable velocity of incoming threats.
- Relative importance (to the enemy) of avoiding detection versus minimizing flight time.
- Level of sophistication of enemy devices for detecting and locating emissions.
- Expected number of aircraft starting at each point (incoming tracks only).
- Probability distribution for the target location.

Outputs

Both models to be described in this section produce displays of the probability of first detection at each point on the grid. In addition, an intermediate output of the first model is a measure of the relative suboptimality of passing through any point on the way to each possible target point T. This quantity will henceforth be denoted the "ratio of suboptimality".

Data Sources

The detection rate is calculated from NTDS information on the status and location of sensor devices and environmental conditions. The perceived detection rate for an observer with imperfect information about the location of Navy sensor devices is computed as described in the section of this appendix devoted to the dynamic programming algorithm. This calculation requires an input parameter describing the sophistication of enemy devices for detecting and locating emissions. The source of this parameter is a subjective assessment by Navy personnel, based on intelligence information.

The relative values of avoiding detection and minimizing flight time, and the maximum sustainable velocity for enemy aircraft, must also be assessed by Navy personnel based on intelligence information.

The target distribution for incoming threats can be estimated from the relative importance of each Navy vessel. Displays have been generated by the experimental decision aids using a number of different target distributions. However, for simplicity the examples in this report assume that the target of each threat is a single ship.

Ratio of Suboptimality Approach

For each target point T, define the ratio of suboptimality $R_T(x)$ by:

$$R_T(x) = \frac{V^*(x) \cdot V_T^*(x)}{V^*(T)},$$

where $V^*(x)$ is the objective value associated with the *optimal path* for reaching point x from the starting set, and $V_T^*(x)$ is the objective value associated with the optimal path from x to T. Thus $R_T(x)$ is the ratio of the value of the best path to point T that passes through point x to the optimal objective value of a path to point T that is not constrained to pass through x.

$R_T(x)$ can be calculated for all x by executing the DP algorithm twice. The first run of the algorithm is used to compute $V^*(i)$ for all i (including, of course, the target point T). On the second run the starting set contains only the point T, with initial condition $V_0(T) = 1$. From this data, the DP algorithm will produce the optimal paths from every point to the target T.

If a target distribution is used, only the latter usage of the DP algorithm must be repeated for each target point receiving

non-negligible weight. Thus a maximum of N+1 executions of the DP algorithm will be required to compute the ratio of suboptimality for all N grid points.

Examples

Consider the simple nine point example of Figure A.5. Let the target point be point 9 (the center point), and let all the other parameters of the problem remain the same. Then $R_9(x)$ has the following values:

Point	Ratio
1.	.8
2.	.8
3.	.51
4.	.51
5.	.41
6.	.64
7.	.64
8.	1.0
9.	1.0

Since the points 3, 4, and 5 are not on or near a good path from the starting set (points 1 and 8) to point 9, they have the lowest ratios. This indicates that it is unlikely that a path passing through any of these points would be chosen.

Figure 8 in the main report gives the display of the relative suboptimality of paths for a more complex example. In this example, the center ship is designated as the target point, attack is allowed from any direction, and the enemy's objective is to maximize his probability of avoiding detection.

Application to First Detection Probability

In order to calculate the first detection probability at each point, we need to know the likelihood that the flight path of an incoming aircraft passes through that point. The ratio of

suboptimality R provides a measure of that likelihood by specifying the relative loss in objective value incurred when the flight path must include that point.

The first detection probabilities are based on the assumption that the probability of passing through a point is proportional to the ratio of suboptimality at that point. In this case, the probability of first detection at a point is given by:

$$P_d(x) = \left(\frac{R(x)}{\sum_i R(i)} \right) \cdot \text{Prob(detect at x on given path)}$$

where the term on the right denotes the probability that detection occurs while near point x if the aircraft chooses to travel a given path containing point x. That probability is conditional on two things: first, that the plane reaches the preceding point on the path (say, point y) without being detected, and second, that the plane is then detected while attempting to travel to point x. The latter quantity is just $1-p(y,x)$, and the former is the product of the nondetection probabilities for travel between the points of the path that precede x. A simple multiplication of the two quantities yields the desired probability of first detection near x for the chosen path.

We have implemented the above model under the assumption that the route chosen to reach each point is the optimal path determined by the dynamic programming algorithm. The first detection displays shown in the main report were produced in this manner.

An alternative method of finding first detection probabilities is to assume the incoming threat will choose from among a set of "reasonable" paths rather than necessarily choosing an optimal path. For a given target point, the ratio of suboptimality algorithm associates a reasonable path with each point x on the grid. This

path connects the starting set with point x, and then continues to the target point T. In order to determine the probability of first detection, we assume the probability of choosing this path is proportional to the ratio $R_T(x)$. Paths or portions of paths that are optimal or near-optimal will have a high likelihood of being chosen, both because the ratio will be higher for these paths, and because they will be repeated many times. For example, the optimal path will, at the very least, be counted once for each point through which it passes.

As before, we can easily determine the probability that detection will first occur at any point, given that one of these paths is chosen. Weighting each of the N paths by the corresponding value $R_T(x)$, we can determine the overall first detection probabilities for the case of a single target by the formula given below.

$$P_d(x) = \sum_{j=1}^{N} \left(\frac{R_T(j)}{\sum_i R_T(i)} \right) \text{Prob(detect at } x \text{ on path } j)$$

To generalize this procedure to multiple targets or a target distribution, simply weight the first detection probabilities calculated individually for each target point by the probability that the point is actually the target.

This procedure should generate first detection displays that are similar in general appearance to the examples shown in the main report. Although this approach poses no computational difficulties, it has not been implemented.

Markov Chain Approach

The main difference between the assumptions necessary for the DP algorithm and the first detection algorithms is that the latter do not make the implicit assumption that an optimum path is always chosen by incoming attackers. In particular, the ratio of suboptimality algorithm assumes that the enemy pilot chooses from a relatively small set of paths, according to their relative merits. Alternatively, the Markov chain model is formulated so that, at each point, the path of an incoming threat is uncertain. In this way, an approaching aircraft can choose to travel along any piecewise linear path through points on the grid, although paths with poor objective values will be chosen infrequently.

The output of this model is the probabilistic location of first detection for incoming aircraft. Alternatively, if an estimate of the size of the next attack is available, the output can be the expected number of planes to be first detected at each location. In addition, the methodology can be extended easily to predict the first detection probabilities for lost tracks.

Formulation

A Markov chain is defined by a set of states, an initial probability distribution over those states, and transition probabilities for moving from one state to another. The initial distribution represents the probability that, at time zero, a quantity (i.e., a threat) is in each possible state. A state is said to be "absorbing" if it is not possible to make a transition from that state to any other state. The probability that a transition to that absorbing state will eventually occur, and hence cause the chain to terminate there, is called the absorption probability.

The movement of a plane within the task force vicinity can be modeled by a Markov chain with the state defined by two factors: the location of the plane and whether or not it has been detected. Since the number of possible locations on the grid is finite, so is the state space of the Markov chain. The one-step transition probabilities of this chain are the probabilities that the plane will choose to travel from a point x to each of the eight neighboring points. A successful transition will occur if the plane is not detected while flying this route, otherwise the plane is absorbed into the "detected" state associated with point x.

Although the standard description of a Markov chain includes an initial probability distribution over the set of states, it is equivalent to allow the initial vector describing the occupation of states to be any set of non-negative numbers. In our case, this leads to a simpler interpretation of the algorithm, in which the initial vector represents the expected number of planes with attack routes starting at each point. To return to the probabilistic formulation, one could divide the expected number of planes at each point by the total expected attack size.

The Markov chain approach requires the following definitions:

$\Pi_0(x)$ = expected number of planes to start at location x (in general these planes can be part of a future attack or lost tracks);

$p(x,y)$ = probability of traveling in a straight line from grid point x to a neighboring grid point y, without being detected;

c = multiplicative reduction in nondetection probability for each additional horizontal or vertical step taken along a path;

$P(y,T)$ = probability of traveling undetected along an "optimal" path from y to target T;

$Z(x,y,T) = c(n_1 + \sqrt{2}n_2) \cdot p(x,y) \cdot P(y,T)$, where n_1 and n_2 represent the number of straight and diagonal steps along the DP-optimal path from x to y to T.

The calculation of $p(x,y)$, the probability of avoiding detection while traveling between two neighboring points, is based on the sensor detection rate, as described in the first section of this appendix. The "optimal" path used in the definitions of the path detection probability $P(y,T)$ and the objective function Z is the path found by the dynamic programming algorithm. Note that $Z(x,y,T)$ is the objective value associated with going from x to y and then taking the DP-optimal path to the target T.

Let each of the N points of the rectangular grid describing the task force vicinity correspond to two states in the Markov chain, one in which the plane is undetected, and an absorbing state in which the plane has been detected. The chain starts in the initial position described by the vector Π_0, and then transitions are allowed between neighbors. The objective function Z is used to determine the transition probabilities Q for this Markov chain. In particular, one alternative is to make the probability of heading from a point x to a neighbor y en route to target T, denoted $Q(x,y,T)$, proportional to the objective function $Z(x,y,T)$.

During an attempted transition from x to y, a plane would be detected with probability $p(x,y)$, and hence absorbed into the appropriate state corresponding to point x. The goal of the algorithm is determine how many absorptions occur at each point, i.e., the expected number of planes to first be detected at that point.

Recall that $Q(x,y,T)$ is the transition probability from x to neighbor y if the target is T. Using the given probability

distribution over the set of targets, we can calculate the overall probability of heading from x to y by:

$$Q(x,y) = \sum_T Q(x,y,T) \cdot \text{Prob(target is T)}.$$

Implementation

To calculate the expected number of planes to be first detected at a point x, we could find the absorption probabilities for this transient chain. For a large number of grid points this calculation might require considerable computer resources; a more efficient technique is described below.

Set Π_0 to the initial distribution of expected planes. Set $E_0(x) = 0$ for all grid points x. Then for an iteration counter $k = 1,2,\ldots,$

$$\Pi_k(x) = \sum_y (Q(y,x)\ p(y,x) \cdot \Pi_{k-1}(y))$$

$$E_k(x) = E_{k-1}(x) + \sum_y (Q(y,x) \cdot (1-p(y,x)) \cdot \Pi_{k-1}(y))$$

Halt when $\sum_x \Pi_k(x) < \varepsilon$, for some positive ε. The parameter ε is the tolerance of the algorithm, and can be set to any desired positive value.

Thus, planes are accumulated in the vector E as they are detected along their flight paths, and planes yet to be detected at iteration k are distributed according to Π_k. As k tends to infinity, E_k converges to the vector of expected number of planes to be first detected at each point. Therefore, if termination occurs at iteration j,

$$E(x) = \frac{E_j(x)}{\sum_i \Pi_0(i)}$$

is an approximation to the first detection probability at x. In fact, it serves as a lower bound, and the total sum of the differences between $E(x)$ and the true absorption probabilities for the Markov chain is less than the tolerance factor ε divided by the total number of planes.

Simple Markov Chain Example

Consider the simple 2x2 grid of Figure A.6. For this example, let a single attacking aircraft begin at point 1, and attempt to reach its target at point 4. A single radar is located at point 4 and surrounded by a symmetric detection field determined by:

$$\begin{aligned} p(1,2) &= .7 = p(1,3) \\ p(2,3) &= .4 \\ p(2,4) &= .3 = p(3,4) \\ p(1,4) &= .1. \end{aligned}$$

For simplicity, let $c=1$, so that the minimum detection paths shown in the figure are optimal.

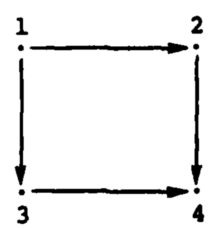

Figure A.6: Minimum Detection Paths for a Simple 2x2 Example

The objective function Z is symmetric with respect to points 2 and 3, and hence is determined by the following values.

$$
\begin{aligned}
Z(1,2) &= .21 = Z(1,3) \\
Z(1,4) &= .1 \\
Z(2,1) &= .147 = Z(3,1) \\
Z(2,3) &= .12 = Z(3,2) \\
Z(2,4) &= .3 = Z(3,4) \\
Z(4,4) &= 1.0
\end{aligned}
$$

The transition probability $Q(1,2)$ is determined by:

$$Q(1,2) = \frac{Z(1,2)}{Z(1,2) + Z(1,3) + Z(1,4)} = \frac{.21}{.51}$$

or approximately .4. Since the probability of avoiding detection along the path from point 1 to point 2 is only .7, the transition from 1 to 2 will only be completed with probability .28.

The complete set of transition probabilities for this Markov chain, including the probability of absorption during attempted transitions, is given in Table A.1 below. Since the target state, point 4, is an absorbing state, it is immaterial whether or not detection is assumed to occur there.

		Transition To				
		1	2	3	4	Detect
Transition From	1	0	.28	.28	.02	.42
	2	.18	0	.08	.16	.58
	3	.18	.08	0	.16	.58
	4	0	0	0	0	1

Table A.1: Transition Probabilities, Including Detection

The first detection probabilities can be found by solving a small system of equations for the absorption probabilities of the Markov chain. The resulting first detecion probabilities are approximately

$$P_d(1) = .47, \quad P_d(2) = .2, \quad P_d(3) = .2, \quad P_d(4) = .13.$$

Note that the probability of first detection at the target site is less than the probability of survival along a minimum detection path (.13 versus .21). This result is due to the fact that a path other than the minimum detection path will sometimes be chosen by incoming threats.

Characteristics of the First Detection Algorithms

1. The first detection algorithms allow probabilistic rather than deterministic choice of paths. Suboptimal paths receive consideration in proportion to their merit. Thus, this type of approach is more robust than the DP method.

2. Using probabilistic models, known, lost, and undetected planes can be aggregated easily. This is a more difficult process under the dynamic programming method due to its deterministic nature.

3. The first detection algorithms require execution of the dynamic programming algorithm once for each possible target point. Thus, they may be up to an order of magnitude slower than the DP approach.

4. The outputs of the first detection algorithms are required inputs for the calculation of expected damage levels.

4. Location Probabilities and Lost Track Algorithms

The major difference between the lost track algorithms and the other methods described thus far is that more information is available about a lost track then an undetected track. In particular, a lost track was detected and observed for some known period of time. From this information, we can calculate its probable current location. In addition, track history and current location can be used to determine more accurately the probable target of the aircraft, and hence to improve the accuracy of the penetration and first detection displays.

Corresponding to the two algorithms developed in the preceding section, there are two methods for evaluating penetration and first detection probabilities for a lost track. This section will concentrate on the formulations of lost track displays utilizing dynamic programming and the ratio of suboptimality. Most of these algorithms have been implemented, although some of the more complex methods are still under development. The relatively simple extension of the Markov chain approach to the lost track case will be described at the end of the section.

Inputs (For All Lost Track Algorithms)

- Sensor detection rate (both actual and perceived).
- Maximum velocity of the threat.
- The location x where the threat was last observed.
- The time interval t (in minutes) since that last observation.
- Target distribution.

Outputs

- Probability distribution over the current location of the track;
- Penetration probability;
- First detection probabilities.

Data Sources

As before, the sensor detection rate is determined from physical data. The maximum velocity can be estimated from the track history and a probabilistic ID associated with the lost track. The remaining input information can be obtained from track history, with the target distribution calculated using kinematics.

Formulation of the Lost Track Current Location Probabilities

The probability that a lost aircraft is currently located at position y can be calculated using Bayes rule. It is proportional to the product of the probability of traveling undetected from x to y and the prior probability that y was the aircraft's destination. Letting $\bar{D}$ be the event of non-detection for t minutes, then

$$\{y|x,D\} = \frac{\{\bar{D}|x,y\}\ \{y|x\}}{\int_y \{D|x,y\}\ \{y|x\}},$$

where the bracketed quantities denote conditional probabilities.

Prior Distributions for Current Location of a Lost Track

The prior distribution for the current location of the track ($\{y|x\}$) can be determined from track history, a target distribution, and knowledge of the aircraft's maximum speed. Using only the latter information, some possible prior distributions are:

(1) A uniform distribution over all the points that the aircraft could reach in t minutes or less.

(2) A uniform distribution over all the points that the aircraft could reach in t minutes or less along the optimum (or minimum detection) routes determined by the dynamic programming algorithm.

(3) A uniform distribution over the outer envelope of either of the regions defined by (1) or (2).

The first two distributions make only very weak statements about the destination of the track. The last distribution reflects the strong assumption that the track is attempting to travel as far from point x as possible.

The prior distribution must also reflect our state of knowledge about the track's target. After choosing a general form for the distribution, (either (1), (2), or (3)), we perform an updating process using a target distribution derived from the application of kinematics to the track history and/or obtained from intelligence information. During this updating process, $L(x)$, the prior probability of being located at point x, is determined by

$$L(x) = \frac{T(x) \cdot U(x)}{\int_Y T(y) \cdot U(y)},$$

where $T(x)$ is the probability that x is the plane's target and $U(x)$ is the value at x of the chosen uniform distribution. This distribution $L(x)$ is a valid prior distribution for the destination of the plane after t minutes.

Figure 17 in Section III of the main report depicts a target distribution using the information that the plane was heading south when last observed. This distribution would be appropriate if the plane is more likely to continue in the same direction than to radically alter its course. Ordinarily this is true, because the incoming plane does not know that it was previously under observation and currently is unobserved.

A target distribution could also be determined by the placement of emitting Navy ships. For example, if the center ship was perceived as the main target, then the target distribution would resemble a set of concentric circles centered at that ship. In general, a target distribution can use target values for all Navy ships to determine the likelihood that each particular ship is the target of the attacking aircraft. These target values would be assessed by Navy personnel from intelligence information.

Probability of Undetected Travel for a Lost Track

Evaluation of the probability of traveling undetected from x to y in t minutes requires an assumption about the trajectory chosen. We assume that the optimum paths found by the DP algorithm are taken. This assumption is reasonable since these paths minimize detection probability; given that no detection has occurred, these paths are more likely to have been chosen. However, note that this criterion excludes some points that could be reached in time t by straight paths. Therefore, we choose not to consider prior distributions of the form of (1).

The probability of traveling undetected from x to y along the optimal paths must be adjusted to reflect the fact that exactly t minutes have passed. This can be done by decreasing the velocity along the path so that the destination y is reached at the appropriate time. This change in velocity can be interpreted as

selection of a path that deviates slightly from the optimal trajectory, and takes longer to complete.

Current Location Display for a Lost Track

The posterior probability for a track to be located at a given point is obtained from the prior distribution and the probability of traveling undetected between that point and the point at which sensor contact was lost. These distributions are combined according to the Bayesian formula previously mentioned.

An example of the display generated by this set of probabilities is given in Figure 18 in Section III of the main report. For this picture, track history (the track was heading south when sensor contact was lost) was used to determine the prior distribution. The display shows the probability of current location for the aircraft ten minutes after the last sensor contact, which occurred at the location marked by the blue dot.

The display of the probable current location for a lost track will change over time. As the time since the last observation increases, the likelihood that the plane is still in the task force vicinity decreases. The rate of this decrease depends on the original reason why the track became lost. There are at least three ways in which sensor contact with a track can be lost:

- Destruction of the track;
- Evasive maneuvers or jamming tactics employed by the enemy;
- Failure of the tracking system, due to mechanical failure of the sensor, range limitations, or reception difficulties.

If a weapon was fired at the track, there is a known possibility that the threat will be destroyed. This possibility must be included in the prior distribution. Since the threat will not be detected after it is destroyed, application of the Bayesian update formula

will increase the likelihood of destruction as time passes without the occurrence of detection.

If the track was lost for another reason, it can be assumed that it continued along its planned flight path. (However, if evasive maneuvers or jamming tactics were employed, that planned flight path may include a radical course change.) As the length of time since the last observation increases, the likelihood that the threat has flown out of sensor range (or been destroyed), which may be near zero initially, will increase. This happens because $\{\bar{D}|x,y\}$ decreases with t for all points y within sensor range. Figures 15 and 16 in the main report illustrate this effect.

For t large enough, the probability that the lost aircraft is still located in the task force vicinity will become negligible, and that track can be dropped from consideration.

Penetration Probabilities for a Lost Track

Calculating the maximum probability of penetration for a lost track is similar to determining penetration probabilities for an incoming plane. However, a lost track will penetrate deeper while avoiding detection than an incoming plane starting at the same point because the lost track is known to have successfully avoided detection for a given length of time.

The penetration probability is calculated using the DP algorithm with the starting set equal to the set of points that can be reached in t minutes or less along the minimum detection paths. Equivalently, detection is assumed not to occur for the first t minutes of travel from point x; all paths are evaluated by the dynamic programming algorithm under this assumption.

The paths determined by this method may not be the same as the minimal detection routes for incoming paths. A suboptimal path which travels through regions of concentrated sensor activity may become optimal when detection is disallowed for the first t minutes of the path. Thus, this maximum penetration probability is truly a worst-case analysis of the situation.

An alternative approach is to evaluate the penetration probabilities for the paths which are optimum before assuming nondetection for t minutes. This approach reflects the fact that the pilot does not have the information of failure to detect when he chooses his path. The probability of penetration along the optimum paths are then evaluated under the constraint that detection can not occur for the first t minutes. This method can easily be generalized by including a term for minimizing flight time in the dynamic programming objective function.

Although this procedure produces a more realistic probability of penetration than the worst-case display, it still does not utilize the updated probabilities of the current location of the track. That data will be utilized for the first detection display.

First Detection Probabilities for a Lost Track

The first detection probabilities for a lost track can be computed using the ratio of suboptimality (described in an earlier section of this appendix). To illustrate the method, suppose there is a single target point T, and let x be the location of the last sensor contact with the track. Then execution of the following steps will produce the first detection probabilities conditional on the target being point T.

1. Compute $R_T(y)$, the ratio of suboptimality for each point y, where the starting set for the dynamic programming subroutine is the single point x. This procedure associates a path from x to y to T with each point y.

2. Compute the probability of being first detected at each point along the path from x to y to T, conditional on failure to detect for the first t minutes along this path.

3. Let the probability of choosing the path associated with point y be proportional to $R_T(y)$.

4. Aggregate over all these paths to find the first detection probability at every point.

This process can be repeated for all target points and the results weighted by the probabilities associated with the targets. The target distribution will be strongly influenced by the information contained in the track history. However, the failure to detect over a period of time implies that points which could have been reached within time t from the original location x are unlikely to have been the target; the target distribution should be updated to reflect this fact.

The procedure outlined above requires a moderate amount of computational effort. The main source of complexity is the requirement that the ratios of suboptimality be determined for each target point. This implies that the DP algorithm must be executed once for each point with non-negligible target probability. Approximately the same order of magnitude of computational effort is needed to implement the Markov chain algorithm, which also can produce first detection probabilities for lost tracks (see the discussion at the end of this section).

A shortcoming of the procedure described above is that knowledge of the failure to detect for t minutes is not used to change the likelihood that each path was chosen (except for a Bayesian update of

the target distribution). An alternative approach is to utilize the probability of current location to specify a distribution for the end point y of whatever path was taken during the t minutes since the track was lost.

Assume first that the current location of the track is known to be point y. Then the steps of the revised algorithm are:

1. Compute $R_T(z)$, the ratio of suboptimality for each point z, where the starting set for the dynamic programming subroutine is the single point y. This procedure associates a path from y to z to T with each point z.

2. Compute the probability of being first detected at each point along the path from y to z to T.

3. Let the probability of choosing the path associated with point z be proportional to $R_T(z)$.

4. Aggregate over all these paths to find the first detection probability at every point, given current location (after t minutes) y and target point T.

This process could be repeated for all target points and all possible current locations, and the results weighted by the appropriate target and current location distributions. Although this algorithm appears to require a great deal of computation, the dynamic programming algorithm need only be executed once for each point on the grid. However, the current location distribution must also be calculated, and this would increase computation time over the algorithm described earlier in this section.

Application of the Markov Chain Approach to Lost Tracks

An alternative method for evaluating first detection probabilities for a lost track is to use the Markov chain approach developed in the last section. For this algorithm, lost tracks can be treated identically to incoming (unobserved) tracks. That is, the

method will calculate the probability of first detection of a lost track provided it is given a starting distribution for the track. That starting distribution can be determined from the probability of current location; as we described earlier in this section, the current location probabilities can be computed using Bayes rule.

The extra information available about the track's past history can be used to improve the accuracy of the target distribution and/or transition probabilities for this aircraft.

The Markov chain approach is currently in the developmental phase. For this reason, the application of this algorithm to lost track first detection probabilities has not yet been implemented.

DISTRIBUTION LIST

OSD

CDR Paul R. Chatelier
Office of the Deputy Under Secretary
of Defense
OUSDRE (E&LS)
Pentagon, Room 3D129
Washington, D.C. 20301

Dr. Stuart Starr
Office of the Assistant Secretary
of Defense (C3I)
Pentagon
Washington, D.C. 20301

Department of the Navy

Director
Engineering Psychology Programs
Code 455
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217 (5 cys)

Director
Aviation & Aerospace Technology
Code 210
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Undersea Technology
Code 220
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Department of the Navy

Director
Communication & Computer Technology
Code 240
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Tactical Development & Evaluation
Support
Code 230
Office of Naval Research
800 North Quincy Street
Arlington VA 22217

Director
Manpower, Personnel and Training
Code 270
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Operations Research Programs
Code 434
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Statistics and Probability Program
Code 436
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Department of the Navy

Director
Electronics & Electromagnetics
Technology
Code 250
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Director
Physiology Program
Code 441
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Special Assistant for Marine
Corps Matters
Code 100M
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Commanding Officer
ONR Eastern/Central Regional Office
ATTN: Dr. J. Lester
Building 114, Section D
666 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
ATTN: Dr. C. Davis
536 South Clark Street
Chicago, IL 60605

Commanding Officer
ONR Western Regional Office
ATTN: Mr. R. Lawson
1030 East Green Street
Pasadena, CA 91106

Commanding Officer
ONR Western Regional Office
ATTN: Dr. E. Gloye
1030 East Green Street
Pasadena, CA 91106

Director
Information Systems Program
Code 437
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Code 430B
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217

Dr. Michael Melich
Communications Sciences Division
Code 7500
Naval Research Laboratory
Washington, D.C. 20375

Dr. Robert G. Smith
Office of the Chief of Naval
Operations, OP-987H
Personnel Logistics Plans
Washington, D.C. 20350

CDR G. Worthington
Office of the Chief of Naval
Operations, OP-372G
Washington, D.C. 20350

Dr. W. Mehuron
Office of the Chief of Naval
Operations, OP 987
Washington, D.C. 20350

Dr. Andreas B. Rechnitzer
Office of the Chief of Naval
Operations, OP 952F
Naval Oceanography Division
Washington, D.C. 20350

Department of the Navy

Office of Naval Research
Scientific Liaison Group
American Embassy, Room A-407
APO San Francisco, CA 96503

Director
Naval Research Laboratory
Technical Information Division
Code 2627
Washington, D.C. 20375 (6 cys)

Human Factors Department
Code N215
Naval Training Equipment Center
Orlando, FL 32813

Dr. Albert Colella
Combat Control Systems
Naval Underwater Systems Center
Newport, RI 02840

Dean of Research Administration
Naval Postgraduate School
Monterey, CA 93940

Mr. Paul Heckman
Naval Ocean Systems Center
San Diego, CA 92152

Dr. Robert French
Naval Ocean Systems Center
San Diego, CA 92152

CAPT Richard L. Martin, USN
Commanding Officer
USS Carl Vinson (CVN-70)
Newport News Shipbuilding/Dry Dock Co.
Newport News, VA 23607

HQS, U.S. Marine Corps
ATTN: CCA40 (MAJOR Pennell)
Washington, D.C. 20380

Dr. Jerry C. Lamb
Combat Control Systems
Naval Underwater Systems Center
Newport, RI 02840

Naval Training Equipment Center
ATTN: Technical Library
Orlando, FL 32813

Dr. Alfred F. Smode
Training Analysis and Evaluation Group
Naval Training Equipment Center
Code N-OOT
Orlando, FL 32813

Dr. Gary Poock
Operations Research Department
Naval Postgraduate School
Monterey, CA 93940

Mr. H. Talkington
Ocean Engineering Department
Naval Ocean Systems Center
San Diego, CA 92152

Mr Warren Lewis
Human Engineering Branch
Code 8231
Naval Ocean Systems Center
San Diego, CA 92152

Dr. Ross L. Pepper
Naval Ocean Systems Center
Hawaii Laboratory
P.O. Box 997
Kailua, HI 96734

Dr. A. L. Slafkosky
Scientific Advisor
Commandant of the Marine Corps
Code RD-1
Washington, D.C. 20380

Commanding Officer
MCTSSA
Marine Corps Base
Camp Pendleton, CA 92055

<u>Department of the Navy</u>

Chief, C^3 Division
Development Center
MCDEC
Quantico, VA 22134

Mr. Arnold Rubinstein
Naval Material Command
NAVMAT 0722 - Rm. 508
800 North Quincy Street
Arlington, VA 22217

Commander
Naval Air Systems Command
Human Factors Program
NAVAIR 340F
Washington, D.C. 20361

Commander
Naval Air Sytems Command
Crew Station Design
NAVAIR 5313
Washington, D.C. 20361

Mr. Phillip Andrews
Naval Sea Systems Command
NAVSEA 0341
Washington, D.C. 20362

Commander
Naval Electronics Systems Command
Human Factors Engineering Branch
Code 4701
Washington, D.C. 20360

LCOL B. Hastings
Marine Corps Liaison Officer
Naval Coastal Systems Center
Panama City, FL 32407

Mr. Milon Essoglou
Naval Facilities Engineering Command
R&D Plans and Programs
Code 03T
Hoffma Building II
Alexandria, VA 22332

CDR Robert Biersner
Naval Medical R&D Command
Code 44
Naval Medical Center
Bethesda, MD 20014

Dr. Arthur Bachrach
Behavioral Sciences Department
Naval Medical Research Institute
Bethesda, MD 20014

CDR Thomas Berghage
Naval Health Research Center
San Diego, CA 92152

Dr. George Moeller
Human Factors Engineering Branch
Submarine Medical Research Lab
Naval Submarine Base
Groton, CT 06340

Head
Aerospace Psychology Department
Code L5
Naval Aerospace Medical Research Lab
Pensacola, FL 32508

Commanding Officer
Naval Health Research Center
San Diego, CA 92152

Dr. James McGrath, Code 302
Navy Personnel Research and
Development Center
San Diego, CA 92152

Navy Personnel Research and
Development Center
Planning & Appraisal
Code 04
San Diego, CA 92152

Department of the Navy

Navy Personnel Research and
Development Center
Management Systems, Code 303
San Diego CA 92152

Navy Personnel Research and
Development Center
Performance Measurement &
Enhancement
Code 309
San Diego, CA 92152

CDR Norman Lane
Code 6021
Naval Air Development Center
Warminster, PA 18974

Dr. Julie Hopson
Human Factors Engineering Division
Naval Air Development Center
Warminster, PA 18974

Mr. Jeffrey Grossman
Human Factors Branch
Code 3152
Naval Weapons Center
China Lake, CA 93555

Human Factors Engineering Branch
Code 1226
Pacific Missile Test Center
Point Mugu, CA ~042

Mr. J. Williams
Department of Environmental
Sciences
U.S. Naval Academy
Annapolis, MD 21402

Dean of the Academic Departments
U.S. Naval Academy
Annapolis, MD 21402

Human Factors Section
Systems Engineerin~ Test
Directorate
U.S. Naval Air Test Center
Patuxent River, MD 20670

Human Factor Engineering Branch
Naval Ship Research and Development
Center, Annapolis Division
Annapolis, MD 21402

Mr. John Quirk
Naval Coastal Systems Laboratory
Code 712
Panama City, FL 32401

CDR W. Moroney
Code 55MP
Naval Postgraduate School
Monterey, CA 93940

Mr. Merlin Malehorn
Office of the Chief of Naval
Operations (OP-115)
Washington, D.C. 20350

Department of the Army

Mr. J. Barber
HQS, Department of the Army
DAPE-MBR
Washington, D.C. 20310

Dr. Joseph Zeidner
Technical Director
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Department of the Army

Director, Organizations and
Systems Research Laboratory
U.S. Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Technical Director
U.S. Army Human Engineering Labs
Aberdeen Proving Ground, MD 21005

U.S. Army Medical R&D Command
ATTN: CPT Gerald P. Krueger
Ft. Detrick, MD 21701

ARI Field Unit-USAREUR
ATTN: Library
C/O ODCSPER
HQ USAREUR & 7th Army
APO New York,, 09403

Department of the Air Force

U.S. Air Force Office of Scientific
Research
Life Sciences Directorate, NL
Bolling Air Force Base
Washington, D.C. 20332

Chief, Systems Engineering Branch
Human Engineering Division
USAF AMRL/HES
Wright-Patterson AFB, OH 45433

Air University Library
Maxwell Air Force Base, AL 36112

Dr. Earl Alluisi
Chief Scientist
AFHRL/CCN
Brooks AFB, TX 78235

Foreign Addresses

North East London Polytechnic
The Charles Myers Library
Livingstone Road
Stratford
London E15 2LJ
ENGLAND

Professor Dr. Carl Graf Hoyos
Institute for Psychology
Technical University
8000 Munich
Arcisstr 21
FEDERAL REPUBLIC OF GERMANY

Dr. Kenneth Gardner
Applied Psychology Unit
Admiralty Marine Technology
Establishment
Teddington, Middlesex TW11 OLN
ENGLAND

Director, Human Factors Wing
Defense & Civil Institute of
Environmental Medicine
Post Office Box 2000
Downsview, Ontario M3M 3B9
CANADA

Dr. A.D. Baddeley
Director, Applied Psychology Unit
Medical Research Council
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND

Other Government Agencies

Defense Technical Information Center
Cameron Station, Bldg. 5
Alexandria, VA 22314 (12 cys)

Dr. Craig Fields
Director, Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Dr. Judith Daly
Cybernetics Technology Office
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Dr. Lloyd Hitchcock
Federal Aviation Administration
ACT 200
Atlantic City Airport, NY 08405

Dr. M. Montemerlo
Human Factors & Simulation Technology, RTE-6
NASA HQS
Washington, D.C. 20546

Dr. J. Miller
Florida Institute of Oceanography
University of South Florida
St. Petersburg, FL 33701

Professor Douglas E. Hunter
Defense Intelligence School
Washington, D.C. 20374

Other Organizations

Dr. Robert R. Mackie
Human Factors Research, Inc.
5775 Dawson Avenue
Goleta, CA 93017

Dr. Gary McClelland
Institute of Behavioral Sciences
University of Colorado
Boulder, CO 80309

Dr. H. McI. Parsons
Human Resources Research Office
300 North Washington Street
Alexandria, VA 22314

Dr. Miley Merkhofer
Stanford Research Institute
Decision Analysis Group
Menlo Park, CA 94025

Dr. Jesse Orlansky
Institute for Defense Analyses
400 Army-Navy Drive
Arlington, VA 22202

Dr. Robert G. Pachella
University of Michigan
Department of Psychology
Human Performance Center
330 Packard Road
Ann Arbor, MI 48104

Professor Judea Pearl
Engineering Systems Department
University of California-Los Angeles
405 Hilgard Avenue
Los Angeles, CA 90024

Professor Howard Raiffa
Graduate School of Business Administration
Harvard University
Soldiers Field Road
Boston, MA 02163

Other Organizations

Dr. T.B. Sheridan
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

Dr. Paul Slovic
Decision Research
1201 Oak Street
Eugene, OR 97401

Dr. Amos Tversky
Department of Psychology
Stanford University
Stanford, CA 94305

Dr. Robert T. Hennessy
NAS - National Research Council
JH #819
2101 Constitution Avenue., N.W.
Washington, D.C. 20418

Dr. Robert Williges
Human Factors Laboratory
Virginia Polytechnical Institute
and State University
130 Whittemore Hall
Blacksburg, VA 24061

Dr. Meredith P. Crawford
American Psychological Association
Office of Educational Affairs
1200 17th Street, N.W.
Washington, D.C. 20036

Dr. Ward Edwards
Director, Social Science Research
Institute
University of Southern California
Los Angeles, CA 90007

Dr. Charles Gettys
Department of Psychology
University of Oklahoma
455 West Lindsey
Norman, OK 73069

Dr. Arthur I. Siegel
Applied Psychological Services, Inc.
404 East Lancaster Street
Wayne, PA 19087

Dr. Harry Snyder
Department of Industrial Engineering
Virginia Polytechnic Institute and
State University
Blacksburg, VA 24061

Dr. W.S. Vaughan
Oceanautics, Inc.
422 6th Street
Annapolis, MD 21403

Dr. M.G. Samet
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. Alphonse Chapanis
Department of Psychology
The John Hopkins University
Charles and 34th Streets
Baltimore, MD 21218

Dr. Elizabeth Kruesi
General Electric Company
Information Systems Programs
1755 Jefferson Davis Highway
Arlington, VA 22202

Dr. Robert Fox
Department of Psychology
Vanderbilt University
Nashville, TN 37240

Dr. Kenneth Hammond
Institute of Behavioral Science
University of Colorado
Room 201
Boulder, CO 80309

Other Organizations

Dr. T.B. Sheridan
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

Dr. Paul Slovic
Decision Research
1201 Oak Street
Eugene, OR 97401

Dr. Amos Tversky
Department of Psychology
Stanford University
Stanford, CA 94305

Dr. Robert T. Hennessy
NAS - National Research Council
JH #819
2101 Constitution Avenue., N.W.
Washington, D.C. 20418

Dr. Robert Williges
Human Factors Laboratory
Virginia Polytechnical Institute
and State University
130 Whittemore Hall
Blacksburg, VA 24061

Dr. Meredith P. Crawford
American Psychological Association
Office of Educational Affairs
1200 17th Street, N.W.
Washington, D.C. 20036

Dr. Ward Edwards
Director, Social Science Research
Institute
University of Southern California
Los Angeles, CA 90007

Dr. Charles Gettys
Department of Psychology
University of Oklahoma
455 West Lindsey
Norman, OK 73069

Dr. Arthur I. Siegel
Applied Psychological Services, Inc.
404 East Lancaster Street
Wayne, PA 19087

Dr. Harry Snyder
Department of Industrial Engineering
Virginia Polytechnic Institute and
State University
Blacksburg, VA 24061

Dr. W.S. Vaughan
Oceanautics, Inc.
422 6th Street
Annapolis, MD 21403

Dr. M.G. Samet
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, CA 91364

Dr. Alphonse Chapanis
Department of Psychology
The John Hopkins University
Charles and 34th Streets
Baltimore, MD 21218

Dr. Elizabeth Kruesi
General Electric Company
Information Systems Programs
1755 Jefferson Davis Highway
Arlington, VA 22202

Dr. Robert Fox
Department of Psychology
Vanderbilt University
Nashville, TN 37240

Dr. Kenneth Hammond
Institute of Behavioral Science
University of Colorado
Room 201
Boulder, CO 80309

Other Organizations

Dr. James H. Howard, Jr.
Department of Psychology
Catholic University
Washington, D.C. 20064

Dr. William Howell
Department of Psychology
Rice University
Houston, TX 77001

Journal Supplement Abstract Service
American Psychological Association
1200 17th Street, N.W.
Washington, D.C. 20036 (3 cys)

Dr. Christopher Wickens
University of Illinois
Department of Psychology
Urbana, IL 61801

Mr. Edward M.Connelly
Performance Measurement
Associates, Inc.
410 Pine Street, S.E.
Suite 300
Vienna, VA 22180

Dr. Thomas P. Piantanida
SRI International
BioEngineering Research Center
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Edward R. Jones
Chief, Human Factors Engineering
McDonnell-Douglas Astronautics
Company
St. Louis Division
Box 516
St. Louis, MO 63166

Dr. Babur M. Pulat
Department of Industrial Engineering
North Carolina A&T State University
Greensboro, NC 27411

Dr. Richard W. Pew
Information Sciences Division
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Hillel Einhorn
University of Chicago
Graduate School of Business
1101 East 58th Street
Chicago, IL 60637

Mr. Tim Gilbert
The MITRE Corporation
1820 Dolly Madison Boulevard
McLean, VA 22102

Dr. David J. Getty
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Douglas Towne
University of Southern California
Behavioral Technology Laboratory
3716 S. Hope Street
Los Angeles, CA 90007

Dr. John Payne
Duke University
Graduate School of Business
Administration
Durham, NC 27706

Dr. Baruch Fischhoff
Decision Research
1201 Oak Street
Eugene, OR 97401

Dr. Andrew P. Sage
University of Virginia
School of Engineering and
Applied Science
Charlottesville, VA 22901

Other Organizations

Dr. Leonard Adelman
Decisions and Designs, Inc.
8400 Westpark Drive, Suite 600
P.O. Box 907
McLean, VA 22101

Dr. A.K. Bejczy
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91125

Mr. Joseph G. Wohl
MITRE Corporation
Box 208
Bedford, MA 01730

Mr. Wayne Zachary
Analytics, Inc.
2500 Maryland Road
Willow Grove, PA 19090

Dr. Lola Lopes
Department of Psychology
University of Wisconsin
Madison, WI 53706

Dr. Stanley N. Roscoe
New Mexico State University
Box 5095
Las Cruces, NM 88003

Dr. Rex Brown
Decision Science Consortium
Suite 721
7700 Leesburg Pike
Falls Church, VA 22043